A Cold Start: Recession Entry and the Purposes of Household Saving*

Tai Lo Yeung†

August 2026

Abstract

Do labor-market conditions at school leaving predict what households save for later in life? I link 4,120 Japanese women to job-openings ratios at predicted graduation. Among component-complete observations when women are married, worse entry conditions are followed by more saving with zeros retained. A direct contrast places most of the estimated response in family-common, child, and other rather than spouse-designated flows. Among positive-saving observations, spouse shares fall, while the differential wife–husband response is imprecise. Estimates use 39 national cohorts and are sensitive to persistent-process inference; within-year evidence is mixed. The results extend labor-market scarring to the purposes of household saving.

Keywords: household finance, labor-market entry, precautionary saving, saving purposes, Japan

JEL: D14, E21, E24, J12, J31

*The JPSC and JHPS/KHPS data were provided by the Panel Data Research Center at Keio University; the restricted prefecture file is used under its access terms. Replication code accompanies the paper; the survey data require application to the PDRC. An online appendix contains supplementary tables.

†Università della Svizzera italiana (USI Lugano). Email: tai.lo.yeung@usi.ch.

1 Introduction

Entering the labor market in a recession can change a career for years. Earnings fall, regular employment becomes harder to obtain, family formation shifts, and asset accumulation is delayed (Kahn, 2010; Oreopoulos et al., 2012; Schwandt and von Wachter, 2019; Genda et al., 2010; Choi et al., 2020). Much less is known about how the household formed after that entry shock responds. A household can insure a scarred member by changing how much it saves, what it postpones, and what the saving is for. Standard balance-sheet data reveal the first margin but rarely the last: deposits and securities are normally recorded for the household as a whole.

This paper studies that missing margin in Japan. The Japanese Panel Survey of Consumers (JPSC) follows 4,120 women over 29 annual waves and asks married respondents how much their household saved in September. It then divides the flow among five designations: family-common, wife, husband, child, and other. I link each woman to the national job-openings-to-applicants ratio in the year she was predicted to leave school. The resulting sample spans 39 graduation cohorts, from 1974 to 2013, including Japan's Employment Ice Age.

The primitive result is an amount response among component-complete observations when women are married. The preferred specification gives every woman total weight one and retains zero saving. A one-unit improvement in the entry ratio predicts 9,900 yen less monthly saving across the five reported components. The non-spouse-designated amount—family-common, child, and other saving—falls by 8,600 yen, while wife- and husband-designated amounts together fall by 1,300 yen. A direct contrast between those aggregates is $-7,250$ yen (cluster $p = 0.012$; score-HAC $p = 0.029$) and is the most extreme of the 39 circular rotations ($p = 1/39 = 0.026$). The point estimates place 87 percent of the total response outside the two spouse categories. In separate outcome samples, adverse entry is also associated with fewer durable purchases and more discouraged borrowing, a pattern consistent with a defensive financial posture.

The familiar share result follows from those amounts. Among positive-saving observations, better entry conditions raise the wife-designated share by 2.2 percentage points and the combined spouse-designated share by 3.6 points on a common five-component sample. The wife amount and the wife–husband split are both estimated imprecisely. The share therefore does not identify a transfer to or from the wife: it is largely the accounting consequence of non-spouse saving expanding more than spouse-designated saving after adverse entry. This distinction recasts the economic question from who controls a fixed pot to what the household saves for.

The designation categories are informative but limited. The question records saving "for" a person or purpose, not legal title, and includes insurance premiums. Wife-designated flows predict growth in the wife's individually held financial assets more strongly than family-common or husband-designated flows, and they predict whether she holds any individual financial asset. These relationships provide construct-validity evidence for designation as a marker of individual attribution. They do not turn it into ownership, and the treatment estimate for the wife's individually held asset stock is imprecise. I consequently make no bargaining-weight or welfare claim.

Identification is the paper's central limitation. The national treatment is one persistent time series observed across 39 predicted graduation cohorts. Age and survey-wave effects absorb a linear cohort trend, but a nonlinear female-cohort process can still mimic the treatment. In the equal-person specification, the direct non-spouse-minus-spouse contrast has a cohort-score HAC value of 0.029. By absolute coefficient magnitude it is the most extreme of 39 circular shifts ($p = 1/39 = 0.026$); studentizing each shifted estimate gives $p = 0.077$. The total-saving and wife-share coefficient-rotation values are 0.128 and 0.154. Thus persistence-aware evidence is strongest for where the estimated amount response occurs, not for whether total saving rises. These results support a descriptive reduced form under the paper's cohort-comparability assumption, not an assumption-free experiment. Spouse-cohort and timing sensitivities are consistent with an entry-year reading for the non-spouse amount but are not decisive (Section 5.2).

I use two within-year comparisons to probe that assumption. A restricted-file design assigns never-movers the job-openings ratio in their current prefecture at predicted graduation and includes graduation-year effects. Adding prefecture effects leaves a positive but imprecise wife-share estimate; the aligned first stage provides no evidence of relevance, and local total-saving and child/other responses disagree with the national ones. A second design replaces the national series with education-tier-specific tightness and absorbs graduation-year effects. Its total and non-spouse level estimates are directionally consistent but imprecise, while the direct allocation contrast is centered on the opposite side of zero and also imprecise. These designs cannot validate the national effect, but reporting them establishes the empirical limit clearly.

Research on saving motives has long treated household saving as purpose-specific rather than fully fungible. Using Japanese microdata, Horioka and Watanabe (1997) decompose saving across twelve motives and show the importance of retirement and precautionary saving. The JPSC instead records realized monthly flows designated as family-common, wife, husband, child, and other. I ask whether labor-market conditions predating the observed household predict both the amount and stated purpose of those flows. This connects motivespecific saving to the recession-entry literature, which documents persistent effects on careers, family formation, health, and household assets (Oreopoulos et al., 2012; Schwandt and von Wachter, 2019; Genda et al., 2010; Kondo, 2024; Choi et al., 2020; Kawaguchi and Kondo, 2020), and to work on saving and portfolio responses to employment risk (Basten et al., 2016; Juelsrud and Wold, 2025; Derenoncourt et al., 2024; Malmendier and Shen, 2024). Intrahousehold-allocation research provides the measurement boundary: designation can be informative without identifying ownership, welfare, or a Pareto weight (Lise and Yamada, 2019; Guo and Xie, 2026; Chiappori et al., 2024; Niimi, 2022; Gu et al., 2026; Belloc and Molina, 2026; Belloc et al., 2025; Yang et al., 2025).

The contribution is deliberately bounded. The data reveal a pooled later-life association between entry conditions and household saving purposes; they do not establish whose welfare changes, whether the response is optimal, or a unique mechanism. The value of the five-way system is that it exposes an amount response and its allocation across purposes that a single share would conceal.

2 Setting, data, and measurement

2.1 Japan's entry market

Japanese firms recruit new graduates through a synchronized annual market. Missing the regular-employment intake can have unusually persistent consequences because movement from nonregular to regular employment has historically been limited (Genda et al., 2010; Kondo, 2007; Hamaaki et al., 2013). The effective job-openings-to-applicants ratio fell from roughly 1.4 around 1990 to about 0.5 in 1999 and did not recover durably until the mid-2000s. Cohorts graduating from 1993 through 2005 are conventionally called the Employment Ice Age generation.

The treatment is the national effective job-openings-to-applicants ratio in a woman's predicted graduation year. Predicted graduation equals birth year plus the modal leaving age for the reported schooling tier: 15 for junior high school, 18 for high school, 20 for vocational or junior college, 22 for university, and 24 for graduate school. Prediction avoids conditioning on an actual completion date that could respond to the economy. It introduces measurement error: predicted and reported years coincide for 43 percent of women, and prediction is one year early for 41 percent. Adjacent treatment years correlate 0.84, and assigning the ratio at the predicted year, the following year, or their average produces similar national estimates.

The ratio is a broad Hello-Work measure and excludes the school-referral market for new graduates. I therefore interpret it as a macro indicator of entry tightness, not the literal queue each graduate faced. A university-specific series gives a right-signed but imprecise amount estimate among university women; a high-school referral ratio and a fully channel-matched construction are weak. Section 5 and Online Appendix Table A4 report these limitations rather than selecting among measures by their outcome fit.

2.2 The JPSC sample

The JPSC began in 1993 with 1,500 women aged 24–34 and added refreshment samples in 1997, 2003, 2008, and 2013 (Panel Data Research Center at Keio University, 2023). I use waves 1–29, through 2021. The resulting panel contains 4,120 women and 53,983 person-waves between ages 24 and 62. Retrospective entry modules record graduation and first-job characteristics. Saving-designation outcomes begin in wave 6.

Table 1: The sample by predicted graduation window

Graduation window	1974–86	1987–92 Bubble	1993–98 Ice I	1999–2005 Ice II	2006–13 Post
Women	1,173	782	795	851	519
Person-waves	21,046	11,750	9,507	8,079	3,601
Regular first job	0.88	0.87	0.75	0.65	0.74
Married, share of waves	0.79	0.70	0.63	0.56	0.46
Median financial wealth, 10,000 yen	200	200	130	120	150
Wife-designated share of saving	0.14	0.14	0.11	0.12	0.11

Notes: JPSC waves 1–29, ages 24–62. Wealth is household deposits plus securities, or the respondent's holdings when single. The wife-designated share is defined for married observations with positive September saving. Later cohorts are observed at younger ages; regressions include age and wave fixed effects.

2.3 Saving amounts, purposes, and validation

Married respondents report total household saving in September. The measure includes insurance-premium payments and is divided into amounts designated as family-common, for the wife, for the husband, for children, and for others (Q294A–F). I use amounts in thousands of yen and retain zero saving. Amount results therefore do not condition on positive saving. Share results apply to married observations with positive reported and component-total saving.

The amount sample requires Q294A–F—the reported total and all five components—to be observed and nonnegative and contains 34,738 person-waves for 2,877 women. The share sample additionally requires positive reported and component-total saving; it contains 31,625 person-waves for 2,794 women. The reported total equals the component sum exactly in 98.53 percent of amount observations and 98.62 percent of share observations. I normalize each share by the sum of the five observed components, so the shares add exactly to one. Forty-seven rows report positive total saving but zero components, while 30 report the reverse; the two consistency conditions remove both cases from the share sample.

The wording is designation, not title. The JPSC uses the explicit term for name or title in separate financial-asset, property, and credit-card questions. Table 4 uses those items to assess what designation tracks. The validation regressions are associational: they include age, wave, and education effects but do not use entry conditions as an instrument.

3 An accounting framework and empirical strategy

3.1 Total saving and its purposes

Let September household saving be

$$S_{it} = F_{it} + W_{it} + H_{it} + K_{it} + O_{it}, \tag{1}$$

where the components are family-common, wife-, husband-, child-, and other-designated amounts. An adverse entry condition can change desired total saving through permanent-income risk, employment insecurity, credit constraints, or experience-based beliefs. It can separately change how the household labels that flow.

For the main comparison, define non-spouse-designated saving as $N = F + K + O$ and spouse-designated saving as $P = W + H$, so $S = N + P$. I also define the allocation contrast $D = N - P$. A negative coefficient on D with respect to entry tightness means that adverse entry is associated with a larger expansion of non-spouse than spouse-designated saving. This partition is exhaustive but semantic: "other" need not benefit the family, and spouse designation need not imply legal title. Its advantage is accounting clarity and a direct test of whether most of the amount response occurs outside the two spouse-specific categories.

The decomposition clarifies what a share can establish. For the wife-designated share,

$$\frac{\partial(W/S)}{\partial R} = \frac{1}{S}\frac{\partial W}{\partial R} - \frac{W}{S^2}\frac{\partial S}{\partial R}, \tag{2}$$

where higher R denotes a tighter entry market. A positive share response can reflect a rising wife amount, a falling total around a fixed wife amount, or both. The same identity applies to the combined spouse amount, $W + H$. The amount system must therefore precede any interpretation of the shares.

The results are organized around three empirical statements. First, $\partial S/\partial R < 0$ together with fewer durable purchases and more borrowing distress after adverse entry would be consistent with a defensive financial posture. Second, the amount contrast $D = N - P$ tests whether most of the estimated response lies outside the two spouse categories. Third, the wife–husband split, $W/(W + H)$, distinguishes a differential response between spouses from a common movement of both spouse categories, conditional on positive spouse-designated saving. These are accounting implications and mechanism diagnostics, not a structural model of household bargaining.

3.2 National specification

For outcome y_{it}, the baseline is

$$y_{it} = \beta R_{g(i)} + \alpha_{a(i,t)} + \tau_t + \gamma_{e(i)} + \varepsilon_{it}, \tag{3}$$

where $R_{g(i)}$ is the national ratio at predicted graduation, and α, τ, and γ are age, survey-wave, and education effects. The treatment varies across 39 graduation years. Age and wave effects mechanically absorb a linear cohort trend; the amount estimates are stable under quadratic and cubic cohort trends, while the wife share attenuates modestly (Online Appendix Table A1). Standard errors cluster on predicted graduation year.

The identifying assumption is that, conditional on these controls, nonlinear cohort movements in the outcome are unrelated to the nonlinear path of the ratio. This assumption is stronger than in region-by-cohort designs with graduation-year effects. The national estimate is consequently best viewed as a reduced form whose causal interpretation depends on cohort comparability.

I report several inferential views rather than choose the smallest result. Graduation-year clustering addresses dependence within the 39 treatment cohorts but treats adjacent cohorts as independent. Exact circular rotations preserve the ordering of the treatment series but require rotational exchangeability; the 39-element reference set includes the observed assignment. Cohort-score HAC permits serial correlation in the exact weighted regression scores. None is an actual randomization design.

3.3 Weighting, selection, and within-year diagnostics

The preferred specification weights each woman's eligible observations by the inverse of her number of observations. It preserves age and wave effects while giving every woman total weight one. Person-wave weighting is reported as a sensitivity check because it gives long-tenure respondents more influence. I also examine family-state controls and potential-experience profiles.

Share outcomes apply to realized married, positive-saving observations. Entry conditions do not precisely predict current marriage or marriage by age 35, and the confidence intervals are wide. Within the married amount sample, a one-unit increase in the ratio predicts a 1.5 percentage-point increase in the probability of any component saving (s.e. 1.5 points) and a 1.6-point increase in share-sample eligibility (s.e. 1.5 points). Amounts including zeros are therefore the cleaner primary outcome; share estimates describe allocation within realized saving households.

Two diagnostics absorb national cohort effects. The prefecture design uses never-movers, assigns the current prefecture's ratio at predicted graduation, and includes graduation-year effects; its demanding form also includes prefecture effects. A second design uses education-tier-specific treatment and graduation-year effects. Both have limited residual variation and are reported as diagnostics rather than alternative headline designs.

4 Results

4.1 The entry exposure predicts the first employment track

Figure 1 shows the institutional first link. Regular first employment and the national job-openings ratio fall together through the ice-age cohorts and recover only partially afterward. Table 2 gives the corresponding regressions. A one-unit improvement in the national ratio raises regular first employment by 7.4 percentage points using reported graduation and by 12.1 points using predicted graduation with a quadratic cohort control. The 0.92 peak-to-trough movement in the ratio therefore corresponds to roughly 7–11 percentage points of regular first employment.

The prefecture first-stage estimate in column 4 is 0.169 (s.e. 0.083) on a broader entry sample of 843 never-movers and uses the prefecture ratio at reported graduation. The aligned local design instead dates exposure at predicted graduation. On amount-sample women, its estimates are 0.0045 (s.e. 0.0317) with graduation-year effects and 0.0071 (s.e. 0.0513) after adding prefecture effects; on share-sample women they are 0.0010 (s.e. 0.0336) and 0.0155 (s.e. 0.0642). I use the first-job results to show that the national ratio marks a consequential entry environment, not as an instrument for a binary first job. Conditions at graduation can operate through the entire career, credit access, beliefs, and family formation, so an exclusion restriction through the first-job indicator is not credible.

Table 2: Entry conditions and regular first employment

	Regular first job			
	(1)	(2)	(3)	(4)
Ratio at graduation	0.074*			
	(0.037)			
Ratio at pred. graduation		0.121***		
		(0.041)		
Ice-age graduate (1993–2005)			-0.129***	
			(0.025)	
Ratio at graduation (pref.)				0.169**
				(0.083)
Standard errors	Grad. yr	Pred. grad. yr	Grad. yr	Pref.
Observations	3,783	3,783	3,783	843
Within Adjusted R^2	0.001	0.007	0.008	0.003
Survey cohort fixed effects	✓	✓	✓	
Education fixed effects	✓	✓	✓	✓
Prefecture fixed effects			✓	✓
Graduation yr fixed effects				✓

Treatment definitions: column 1, the national job-openings-to-applicants ratio in the reported graduation year; column 2, the national ratio in the year predicted from birth year and schooling track (15/18/20/22/24 for junior high/high school/vocational or junior college/university/graduate school), with a quadratic in predicted graduation year and age controlled (coefficients suppressed); column 3, an indicator for reported graduation in 1993–2005; column 4, the prefecture-level ratio in the reported graduation year, on never-movers (cohorts B–E), for whom current prefecture equals graduation prefecture. Columns 1–3: full entry sample. Standard errors in parentheses, clustered on graduation year (columns 1–3) and prefecture (column 4); $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

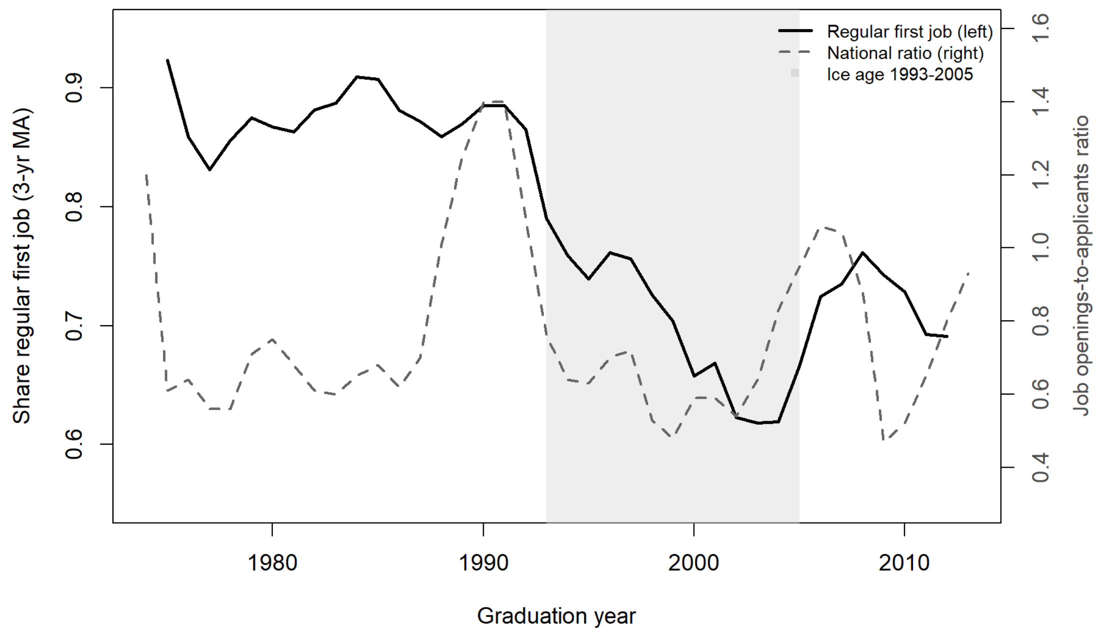

Figure 1: Labor-market tightness and regular first employment. The solid line is the three-year moving average of the share entering a regular first job; the dashed line is the national job-openings-to-applicants ratio. Shading marks predicted graduation in 1993–2005. JPSC entry modules, 3,783 women.

4.2 Saving amounts with zeros retained

Panel A of Table 3 is the paper's central result. Among component-complete observations when women are married, higher entry tightness predicts lower subsequent saving: the coefficient on the sum of the five components is -9.89 thousand yen (cluster s.e. 3.18), relative to an equal-person-weighted mean of 68.76 thousand. The cohort-score HAC standard error is 3.90 ($p = 0.016$), while 5 of 39 circular shifts have coefficients at least as large in absolute value ($p = 0.128$). The total response is therefore clear under conventional and score-HAC inference but not unusual under every persistent treatment path.

The allocation contrast locates that response more sharply. The non-spouse-designated amount—family-common, child, and other saving—has a coefficient of -8.57 thousand yen, compared with -1.32 for spouse-designated saving. Their direct difference, $N - P$, is -7.25 thousand yen (cluster s.e. 2.76; score-HAC s.e. 3.19). Its clustered and score-HAC values are 0.012 and 0.029. By absolute coefficient magnitude it is the most extreme of the 39 circular rotations ($p = 1/39 = 0.026$); a studentized rotation comparison gives $p = 0.077$. Thus the data distinguish the aggregate amount responses under the clustered and score-HAC views, with suggestive rather than conventionally decisive studentized-rotation evidence; the claim does not rely on treating the spouse coefficients as zeros. Family-common saving contributes -6.02 thousand yen, child saving -1.36, and other saving -1.20, so the point estimate is driven primarily by the family-common component. At the 0.48 difference between the late-bubble and trough cohort means, the estimates imply about 4,750 yen more five-component saving after adverse entry, of which roughly 4,110 yen—87 percent—is non-spouse-designated.

The total-saving sign is consistent with a defensive financial posture. In separate national outcome samples, worse entry predicts fewer durable purchases (5.4 percentage points per unit of the ratio; s.e. 2.5) and more discouraged borrowing (2.5 points per unit; s.e. 0.7) alongside more saving (Online Appendix Table A2). These outcomes are supporting evidence, not additional primary outcomes. Wealth ranks, homeownership, securities participation, and the wife's individually held stocks are imprecise, so annual flows should not be translated into a demonstrated long-run wealth effect.

The non-spouse amount is not an artifact of a few very large reports. Its coefficient remains negative when I winsorize each outcome at either the 99th or 95th percentile. The same is true under exact reported-total reconciliation, a common age-30–45 window, and exclusions of the top 1 percent of observations or women. The five-component total is more tail-sensitive: it remains negative and precise under 99th-percentile winsorization and exact reconciliation, is borderline at the 95th percentile ($p = 0.052$), and becomes imprecise after the two top-tail exclusions (Online Appendix Table A7). A two-part description points to the intensive margin. Entry tightness barely predicts whether any component saving is positive, while conditional on a positive component total it predicts 9.76 thousand yen less non-spouse saving (s.e. 2.72) and 11.68 thousand yen less five-component saving (s.e. 3.53). Because entry into the positive-saving subsample can itself be selected, those conditional regressions are descriptive support rather than a decomposition of the fixed-effect coefficient (Online Appendix Table A6).

Table 3: Labor-market entry conditions and the purposes of household saving

Outcome	Mean	$\hat{\beta}$	Cluster s.e.	Score-HAC(5) s.e.	Rotation p	Women
Panel A. Amounts with zeros retained (thousand yen)						
Five-component total	68.76	−9.89	3.18	3.90	0.128	2,877
Non-spouse-designated amount (N)	49.40	−8.57	2.48	3.25	0.051	2,877
Spouse-designated amount (P)	19.36	−1.32	1.64	1.46	0.615	2,877
Allocation contrast ($N - P$)	**30.04**	**−7.25**	**2.76**	**3.19**	**0.026**	**2,877**
Family-common amount	36.69	−6.02	2.19	2.75	0.051	2,877
Wife-designated amount	8.34	−0.25	0.89	0.69	0.872	2,877
Husband-designated amount	11.02	−1.06	0.93	0.85	0.410	2,877
Child-designated amount	10.43	−1.36	0.75	0.89	0.513	2,877
Other-designated amount	2.28	−1.20	0.49	0.36	0.103	2,877
Panel B. Normalized five-component shares (sum exactly to one)						
Family-common share	0.493	−0.015	0.016	0.019	0.487	2,794
Wife-designated share	0.120	0.022	0.009	0.007	0.154	2,794
Husband-designated share	0.166	0.015	0.013	0.012	0.103	2,794
Child-designated share	0.189	−0.012	0.014	0.012	0.513	2,794
Other-designated share	0.032	−0.009	0.006	0.006	0.256	2,794
Non-spouse-designated share	0.714	−0.036	0.018	0.017	0.051	2,794
Panel C. Within spouse-designated saving, conditional on $P > 0$						
Wife share within spouse-designated saving	0.425	−0.005	0.015	0.013	0.744	2,368

Notes: The treatment is the national job-openings-to-applicants ratio in the woman's predicted graduation year. All regressions include age, survey-wave, and education fixed effects and give every woman total weight one within the relevant sample. The amount sample contains married observations with the reported total and all five components observed and nonnegative, including zero saving; the share sample additionally requires positive reported and component-total saving. N is family-common plus child plus other saving; P is wife- plus husband-designated saving. The within-spouse row further requires $P > 0$. Cluster standard errors group by 39 predicted graduation years. Score-HAC standard errors aggregate exact weighted regression scores by ordered graduation year with five Bartlett lags. Rotation p-values compare absolute coefficients over all 39 circular shifts, including the observed assignment; they are dependence-preserving placebos, not randomization inference. For $N - P$, studentizing every shifted coefficient gives $p = 0.077$ rather than the coefficient-magnitude value of 0.026 shown in the table.

4.3 Shares are a numerator–denominator decomposition

Panel B reports the same system in shares. Better entry conditions raise the wife-designated share by 2.15 points and the husband-designated share by 1.45 points; together the spouse share rises by 3.60 points. The exact non-spouse share falls by the same amount because the five normalized shares add to one. The individual family-common, child, and other coefficients are all negative, although none is individually unusual across circular shifts.

Equation (2) explains the pattern. Better entry conditions reduce total and non-spouse saving, while the two spouse amounts have smaller estimated responses. Spouse-designated amounts consequently occupy larger shares of a smaller total. In the adverse-entry direction, non-spouse saving expands more than spouse-designated saving, causing both spouse shares to fall. Conditional on positive spouse-designated saving, the wife–husband split has a coefficient of -0.005 (cluster s.e. 0.015; 95 percent interval approximately $[-0.034, 0.024]$; score-HAC $p = 0.705$; rotation $p = 0.744$). The estimates therefore do not establish a transfer between spouses or a change in a sharing rule; their interval permits moderate differential responses.

The shares are secondary because persistence-aware inference is materially weaker than for the direct amount contrast. The wife-share coefficient has a score-HAC value of 0.004 but a circular-shift value of $6/39 = 0.154$; the non-spouse share has corresponding values 0.044 and $2/39 = 0.051$. An asymptotic score-HAC Wald test rejects that four independent share coefficients are jointly zero ($p = 0.018$). A circular-shift max statistic is sensitive to which adding-up share is omitted, so I do not use it to rank the share system. The share pattern is informative about the accounting decomposition, not an independently secure effect.

4.4 What designation tracks

Table 4 provides construct-validity evidence for the survey label without converting it into legal ownership. The wife-designated share correlates 0.18 with her share of individually reported financial assets, and a one-point higher designation share predicts a 0.23-point higher asset share conditional on age, wave, and education. In flow-to-stock regressions, an additional thousand yen of wife-designated September saving predicts 0.30 units of next-year growth in her individually reported assets, compared with 0.09 for family-common saving and an imprecise -0.01 for husband-designated saving. The presence of any wife-designated saving also predicts any individually held financial asset and a credit card in her name.

These are associational validation results. They show that respondents use the categories meaningfully. They do not establish legal title for Q294, and the entry-treatment estimate for the wife's static individual asset stock is imprecise. The paper therefore describes purpose and attribution, not ownership or welfare.

Table 4: Construct validation: designation and individually reported assets

Outcome and predictor	Coefficient	Standard error	Observations
Wife asset share ← wife-designated share	0.226	0.040	24,888
Growth in wife's assets ← wife amount	0.300	0.118	29,415
Growth in wife's assets ← common amount	0.094	0.040	29,415
Growth in wife's assets ← husband amount	−0.012	0.072	29,415
Any individual financial asset ← any wife designation	0.160	0.011	34,586
Own-name credit card ← any wife designation	0.023	0.008	21,240

Notes: Associational validation regressions, not treatment effects. Models include age, survey-wave, and education fixed effects. Asset growth is the change from the current to the next wave. Amounts are thousand yen in September; stock changes are in 10,000 yen. Standard errors cluster on the respondent. Q294 includes saving and insurance-premium payments and records designation, not title.

5 Credibility, inference, and limits

5.1 Effective sample size and persistent exposure

The regression has tens of thousands of person-waves, but its national treatment support consists of 39 graduation cohorts. Table 3 therefore reports three inferential views for each outcome on its own fixed specification and sample. Graduation-year clustering allows arbitrary dependence within a cohort. Score-HAC aggregates the exact weighted regression scores by ordered cohort and applies five Bartlett lags. Circular shifts keep the treatment series intact and compare the observed coefficient with all 38 rotations; the observed assignment is included, so attainable values are multiples of 1/39. Rotations are dependence-preserving placebos under rotational exchangeability, not randomization inference.

The hierarchy matters. The direct non-spouse-minus-spouse amount contrast is precise under clustering and score-HAC. Its absolute coefficient is the most extreme circular rotation ($p = 0.026$), although the studentized comparison gives $p = 0.077$. Total saving is precise under the first two procedures but not coefficient rotations ($p = 0.128$). The wife share is also precise under the first two procedures but not coefficient rotations ($p = 0.154$). An asymptotic score-HAC test rejects that the five amount coefficients are jointly zero ($p = 0.0025$), whereas their circular-shift max statistic does not ($p = 0.205$). The analogous share-system rotation statistic depends on the omitted adding-up category and is not used for inference. The best-supported result is therefore the relative amount response across the non-spouse and spouse aggregates, not the total or any individual share.

Figure 2 visualizes the direct allocation contrast using the exact weighted Frisch–Waugh–Lovell regression. Both $D = N - P$ and the entry ratio are residualized on the same age, survey-wave, and education effects. The fitted slope is -7.25, numerically identical to the individual weighted fixed-effect coefficient in Table 3; the binned points are only a visual guide. The contrast remains negative in all 39 leave-one-cohort regressions, ranging from -9.08 to -6.19. Dropping both endpoint cohorts gives -7.28 (cluster s.e. 2.77), so the result is not generated by either edge of the treatment support.

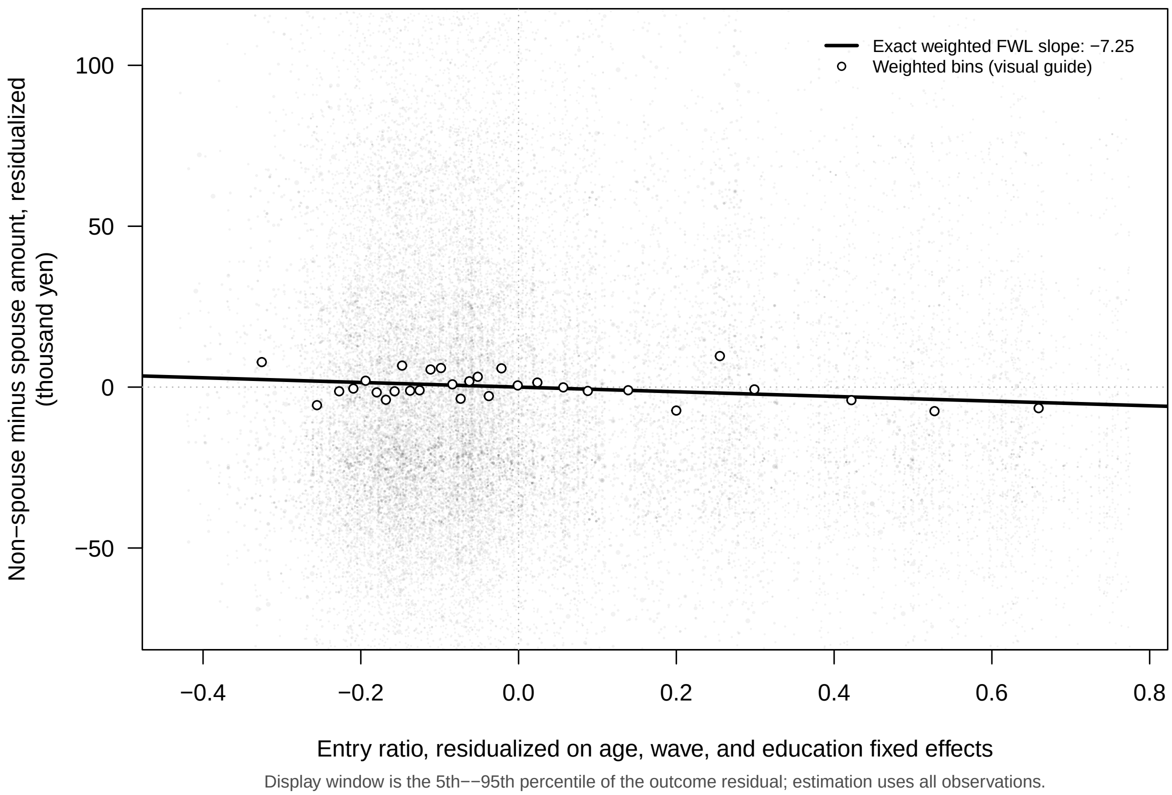

Figure 2: Exact partial regression for the allocation contrast. The horizontal axis is the national entry ratio residualized on age, survey-wave, and education fixed effects; the vertical axis is $D = N - P$ residualized on the same effects. Gray points are person-waves shown within the fifth–95th percentile display window of the outcome residual; open circles are weighted bins. The black line uses every observation and has slope -7.25, exactly reproducing the equal-person-weighted fixed-effect coefficient.

5.2 Spouse-cohort and timing sensitivities

Two comparisons probe the dating and spouse specificity of the national association; neither rules out a nonlinear female-cohort process (Online Appendix Table A3). First, adding the husband's graduation-year ratio and husband-education effects leaves the wife's own-ratio estimate negative for the non-spouse amount and positive for the wife share. A formal own-minus-husband contrast rejects equality for the amount ($t = -2.17$, $p = 0.036$) but not for the wife share ($t = 1.42$, $p = 0.165$). In the amount model (34,669 person-waves; 2,870 women), husband graduation year is reported for 1,444 respondent IDs and predicted for 1,426, which may attenuate his coefficient; moreover, his cohort is a post-treatment match outcome. The comparison is therefore a spouse-cohort sensitivity, not a falsification of female-specific cohort confounding.

Second, the ratio five years after predicted graduation is correlated -0.22 with the own-year ratio across the 39 cohorts, so its opposite-signed standalone amount coefficient is not independent evidence. In the joint regression, the own-year minus $g+5$ contrast is -9.83 thousand yen for the non-spouse amount ($p = 0.0008$), but only 0.015 for the wife share ($p = 0.152$). The joint specification distinguishes the two coefficients for the non-spouse amount but not for the wife share. Because both regressors remain transformations of one national time series, the exercise sharpens dating within the specification without eliminating nonlinear cohort confounding. Parents' education likewise has no detected gradient in the entry ratio ($t = 0.7$; $t = -0.1$ with a quadratic cohort trend), but those estimates do not rule out moderate imbalance.

5.3 Weighting and family-state diagnostics

Equal-person weighting is the baseline because panel tenure differs sharply across refreshment cohorts. Table 5 shows the corresponding person-wave estimates. They are larger in magnitude, showing sensitivity to the additional influence assigned to longer-tenure respondents. Adding child-count and observed-marriage-duration effects leaves the direct allocation contrast negative, although those controls may themselves respond to entry conditions and are diagnostic rather than preferred.

Table 5: Weighting and family-state diagnostics

Outcome	Specification	Coefficient	Cluster s.e.
Allocation contrast, $N-P$	Equal-person baseline	-7.252	2.755
	Person-wave weighting	-10.251	2.256
	Family-state effects	-7.362	2.827
Five-component total	Equal-person baseline	-9.887	3.181
	Person-wave weighting	-14.401	3.316
	Family-state effects	-10.864	3.356
Wife-designated share	Equal-person baseline	0.0215	0.0090
	Person-wave weighting	0.0235	0.0068
	Family-state effects	0.0165	0.0083

Notes: National specifications include age, survey-wave, and education effects; standard errors cluster on 39 predicted graduation years. Equal-person weights are the inverse of a woman's eligible observations. Person-wave rows are unweighted. Family-state rows use equal-person weights and add child-count and observed-marriage-duration effects; those variables may be post-treatment. Amounts are thousand yen. $N-P$ is non-spouse-designated saving (family-common plus child plus other) minus spouse-designated saving (wife plus husband).

5.4 Within-year evidence is mixed, not validation

The local design is most credible with both graduation-year and prefecture effects. Its wife-share estimate is positive but imprecise, 0.025 (s.e. 0.033), and its direct $N-P$ contrast is also imprecise, 5.76 thousand yen (s.e. 10.05; Online Appendix Table A5). On amount-sample women, the first-stage estimates are 0.0045 (s.e. 0.0317) with graduation-year effects and 0.0071 (s.e. 0.0513) after prefecture effects are added; on share-sample women they are 0.0010 (s.e. 0.0336) and 0.0155 (s.e. 0.0642). The broader entry-sample first stage in Table 2 cannot substitute for relevance in the selected outcome sample.

The local amount system also disagrees with the national one. Better local entry conditions predict 23.30 thousand yen more total saving with graduation-year effects and 32.46 thousand more after prefecture effects are added; the corresponding non-spouse estimates are 12.15 and 19.11 thousand. The local child/other share also rises with better entry conditions, opposite the national sign. These differences may reflect limited support, migration-based exposure error, or local prosperity entering amount outcomes. Whatever their source, they prevent the local exercise from being described as independent confirmation.

The education-tier design reaches the same limit from a different direction. Graduation-year and baseline effects leave 26.2 percent of the treatment standard deviation in the raw amount sample and 26.5 percent in the common share sample—about 7 percent of its variance. Per one equal-person-weighted within-tier standard deviation, the raw-sample non-spouse, five-component-total, and direct $N-P$ estimates are -2.22 (s.e. 3.27), -5.22 (s.e.

3.84), and 0.78 thousand yen (s.e. 3.74); the common-sample wife-share estimate is 0.009 (s.e. 0.010). None is precise, and the tier design does not reproduce the national allocation contrast. The route to stronger identification is actual graduation geography or an entrant-specific demand shock with cohort effects and a first stage in the same outcome sample.

5.5 Selection, family state, and horizon

Amounts include reported zeros, but the sample remains conditional on realized marriage. In the equal-person specification, the coefficient on being married is 0.036 (s.e. 0.034), and that on being observed married by age 35 is 0.001 (s.e. 0.056); neither rules out meaningful selection. Among married component-complete observations, the coefficients are 0.015 (s.e. 0.015) for any positive component saving and 0.016 (s.e. 0.015) for share-sample eligibility. Within the common positive-saving sample (31,625 person-waves; 2,794 women), the coefficient on positive spouse-designated saving is 0.0355 (s.e. 0.0220). A supplementary Lee-style bound for the binary ice-age wife-share contrast spans zero (Online Appendix Table A6). Amount and share results should therefore remain explicitly conditional on realized marriage, and shares additionally on positive saving.

Family composition is especially important because children create a designated category and the saving split changes around childbirth. The direct $N-P$ coefficient remains negative in childless and parent observations, but both estimates are imprecise (Online Appendix Table A8). I treat those splits as diagnostics rather than a mechanism. Likewise, the pooled specification does not establish a constant effect for three decades. Allocation-contrast coefficients are negative in all four potential-experience bins, but range from -3.16 at 0–9 years to -15.41 at 30 or more years while cohort support changes. The revised claim is a pooled later-life association, not one constant thirty-year causal scar.

6 Interpretation

The national pattern is compatible with a purpose-specific form of self-insurance, but does not establish it. Adverse entry is associated with more saving, fewer durable purchases, greater discouraged borrowing, and a larger response outside the two spouse-specific categories. One interpretation is that employment insecurity and restricted credit raise the value of a common buffer and planned family expenditure. Other interpretations remain possible: the categories include insurance premiums, permanent-income effects need not have a single sign, and cohort norms can affect how identical financial products are labeled.

Three interpretations are not supported. First, the result does not establish a transfer from wife to husband; both spouse shares move together, while their internal split is centered near zero and imprecise. Second, it is not a demonstrated fall in the wife's individually held wealth; her amount and stock treatment effects are imprecise. Third, it is not a welfare ranking of household and individual saving. Designation predicts individual attribution cross-sectionally, but the survey does not reveal eventual use or utility.

The setting also limits external validity. Japan combines synchronized graduate recruitment, persistent regular/nonregular segmentation, and explicit household-money conventions. A U.S. SIPP exercise using coarser national unemployment exposure and individual asset allocation yields estimates centered near zero (Online Appendix Table A9). Korea's spouse-specific pension data provide a closer next setting (Yang et al., 2025), but a comparable entry design has not yet been implemented.

7 Conclusion

Labor-market entry conditions are usually evaluated by what they do to careers. This paper asks whether they also predict what a household later saves for. In Japan's Employment Ice Age, worse conditions at school leaving were associated with a lower probability of regular first employment. Among component-complete observations when women are married, those conditions are followed by more saving with zeros retained, with most of the estimated additional yen outside the wife- and husband-designated categories. The direct $N - P$ contrast, rather than the insignificance of either spouse amount, supports that allocation claim. Among positive-saving observations, spouse shares fall, while the wife–husband split is imprecisely estimated.

This amount pattern changes the interpretation. It is not evidence that saving is transferred from wife to husband, that the wife's bargaining weight falls, or that her individually held wealth is reduced. It is a difference in estimated purpose-specific responses in the national reduced form: the point estimate is driven primarily by family-common saving, with smaller child and other components. Lower durable purchasing and greater discouraged borrowing are consistent with a defensive financial posture, although the data do not identify the preferences or constraints that generate it.

The strength of that conclusion is bounded by the design. The national exposure is a persistent series observed across 39 effective cohorts. The direct allocation contrast is the most extreme circular shift by coefficient magnitude ($p = 0.026$), although the studentized comparison gives $p = 0.077$; total saving has a coefficient-rotation value of 0.128. Persistence-aware evidence is therefore stronger for where the estimated amount response occurs than for whether total saving rises. Within-year prefecture estimates do not reproduce the amount system, and the aligned local first stage provides no evidence of relevance. Selection into marriage and positive saving and changing family state further limit the estimand. The findings should therefore be read as a transparent national cohort reduced form, not as a fully isolated causal parameter.

Even at that level, the result adds a dimension to labor-market scarring that household totals conceal. Two households can report similar spouse-designated saving while differing substantially in the common and child-related saving built around it. Following unlucky cohorts into the household therefore requires observing both how much families save and the purposes attached to those flows. Whether the pattern represents effective insurance, constrained consumption, or another response requires stronger within-cohort variation and data linking saving purposes to later use.

References

Christoph Basten, Andreas Fagereng, and Kjetil Telle. Saving and portfolio allocation before and after job loss. *Journal of Money, Credit and Banking*, 48(2-3):293–324, 2016.

Ignacio Belloc and José Alberto Molina. How inheritance expectations impact household savings. *Review of Income and Wealth*, 72(1):e70044, 2026. doi: 10.1111/roiw.70044.

Ignacio Belloc, Pierre-André Chiappori, José Alberto Molina, and Jorge Velilla. Effects of wage shocks and saving changes on leisure time: The role of dynamic intra-household commitment. Boston College Working Papers in Economics 1099, 2025.

Pierre-André Chiappori, Costas Meghir, and Yoko Okuyama. Intra-household welfare inequality and household public goods. NBER Working Paper 32645, 2024.

Eleanor Jawon Choi, Jaewoo Choi, and Hyelim Son. The long-term effects of labor market entry in a recession: Evidence from the Asian financial crisis. *Labour Economics*, 67:101926, 2020.

Ellora Derenoncourt, Chi Hyun Kim, Moritz Kuhn, and Moritz Schularick. Unemployment risk, portfolio choice, and the racial wealth gap. Working paper, 2024.

Yuji Genda, Ayako Kondo, and Souichi Ohta. Long-term effects of a recession at labor market entry in Japan and the United States. *Journal of Human Resources*, 45(1):157–196, 2010.

Ran Gu, Cameron Peng, and Weilong Zhang. The gender gap in household bargaining power: A revealed-preference approach. *Review of Financial Studies*, 39(6):1611–1653, 2026.

Naijia Guo and Anning Xie. Childbirth and welfare inequality: The role of bargaining power and intrahousehold allocation. *Journal of Labor Economics*, 2026. Forthcoming.

Junya Hamaaki, Masahiro Hori, Saeko Maeda, and Keiko Murata. How does the first job matter for an individual's career life in Japan? *Journal of the Japanese and International Economies*, 29:154–169, 2013. doi: 10.1016/j.jjie.2013.07.004.

Charles Yuji Horioka and Wakō Watanabe. Why do people save? a micro-analysis of motives for household saving in Japan. *The Economic Journal*, 107(442):537–552, 1997. doi: 10.1111/j.1468-0297.1997.tb00026.x.

Ragnar E. Juelsrud and Ella Getz Wold. The saving and employment effects of higher job loss risk. Working paper, 2025.

Lisa B. Kahn. The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics*, 17(2):303–316, 2010.

Daiji Kawaguchi and Ayako Kondo. The effects of graduating from college during a recession on living standards. *Economic Inquiry*, 58(1):283–293, 2020.

Ayako Kondo. Does the first job really matter? state dependency in employment status in Japan. *Journal of the Japanese and International Economies*, 21(3):379–402, 2007.

Ayako Kondo. Scars of the job market "ice-age". *Social Science Japan Journal*, 27(2):133–148, 2024. doi: 10.1093/ssjj/jyae008.

Jeremy Lise and Ken Yamada. Household sharing and commitment: Evidence from panel data on individual expenditures and time use. *Review of Economic Studies*, 86(5):2184–2219, 2019.

Ulrike Malmendier and Leslie Sheng Shen. Scarred consumption. *American Economic Journal: Macroeconomics*, 16(1):322–355, 2024.

Yoko Niimi. Are married women really wealthier than unmarried women? evidence from Japan. *Demography*, 59(2):461–483, 2022. doi: 10.1215/00703370-9735271.

Philip Oreopoulos, Till von Wachter, and Andrew Heisz. The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1):1–29, 2012.

Panel Data Research Center at Keio University. Japanese panel survey of consumers (JPSC), waves 1–29. Keio University, 2023.

Hannes Schwandt and Till von Wachter. Unlucky cohorts: Estimating the long-term effects of entering the labor market in a recession in large cross-sectional data sets. *Journal of Labor Economics*, 37(S1):S161–S198, 2019.

Hae Kyung Yang, Sharon Tennyson, and Frances Woolley. In whose name? joint planning of individual retirement savings in Korea. *Applied Economics*, 2025. doi: 10.1080/00036846.2025.2593678.