

Impact of population aging on youth wages in Japan: Evidence from Japanese panel data

CSISZÁR Csenge¹

DOI: [10.29180/978-615-6886-26-2_6](https://doi.org/10.29180/978-615-6886-26-2_6)

Abstract

Japan is now a super-aged society, with a large share of the population aged 65 or older. In this situation, the traditional labour structure, where wages are linked to age and tenure under lifetime employment, still plays a strong role. Many companies continue supporting elderly workers assuming their experience brings productivity, and government policies encouraging work beyond retirement age help maintain this system. While official data do not report negative effects of elderly employment on youth wages, international research suggests that delayed retirement can reduce opportunities for younger people.

This study explores whether a similar dynamic is occurring in Japan by using grouped panel data covering employment by age group. An Overall Wage Score was developed to track wage distribution trends across cohorts, and a Bayesian multilevel model was applied to estimate the effects of changes in high-income groups. The model was estimated using the Japanese Panel Study of Employment Dynamics (JPSED). Results suggest that older age groups are gaining wage share while younger groups are losing ground, especially in lower brackets. The model indicates a negative relationship between growth in high-income cohorts and the wage score for others, which supports the idea that structural changes may put younger workers at a disadvantage.

Due to the use of grouped data, important variables like occupation or hours worked could not be included, which limits interpretation. Future studies using individual-level data could provide deeper insight into wage inequality in Japan's aging workforce.

Keywords: Ageing, Labor, Wages, Population, Japan

JEL Classifications: J82

Introduction

Japan is currently going through one of the most advanced demographic transitions globally. The country has already entered a stage called a “super-aged society,” where more than 28% of the population is aged 65 and over. This transformation already shows effects in the structure of the labour market, most notably seen in the flattening of the age–wage profile, where wage increases by age have become less steep than before (Abegglen, 1958). At the same time, the Japanese government is responding to demographic change by actively promoting elderly employment rather than restricting it. According to the 2024 White Paper on Health, Labour and Welfare published by the Ministry of Health, Labour and Welfare (MHLW) (MHLW, 2024), policies now aim to build a society where “everyone can remain active throughout their lives.” Reforms include increasing the retirement age to 65 and encouraging workers to stay employed beyond 70. These policies are meant to fill labour shortages, and the report does not claim that elderly workers cause a decline in youth wages. The report also mentions that younger workers often feel stress caused by their workload and responsibility, but these are not shown as consequences of elderly

¹ Waseda University, csiszar.75@fuji.waseda.jp

employment. However, international evidence shows that there can be indirect effects of delayed retirement on young workers. In the United States, research shows that when older workers retire later, this creates a “bottleneck” effect thus reducing available jobs for youth. This can cause more competition, lower starting wages, and fewer chances for promotion, all of which can affect long-term wage growth for younger people. Even though this is based on U.S. data, a similar effect could possibly happen in Japan too, especially in a labour system where promotions and wages are strongly based on age (Mohnen & Strauss, 2024).

Ghilarducci, Papadopoulos, and Radpour (2017) examined intergenerational labour market dynamics and found that increases in the supply of older workers are associated with downward pressure on wages for younger cohorts in the United States of America. An increase in the relative labour supply of older men is associated with lower wages for young women. A rise in the relative supply of older women is linked with lower wages for young and prime-aged men. This suggests a shift in labour market substitution patterns, where older workers replace younger ones more frequently than in earlier periods (Ghilarducci et al., 2017). Over a ten-year period, the projected addition of 7.2 million workers is expected to impact wage growth. If the composition of age and sex groups in the workforce remains constant, annual wage growth between 2014 and 2024 is estimated at 1.7 percent. However, labour supply forecasts show uneven wage growth across groups. Young women and men are expected to experience wage growth of 1.81 percent and 1.18 percent, respectively, while older women and men are projected to see wage increases of 4.8 percent and 7.48 percent (Ghilarducci et al., 2017), respectively. Older workers’ increased participation corresponds with slower wage growth for younger cohorts. Policy changes that raise the retirement age or extend Medicare eligibility to age 70 are linked with broader labour market effects, including wage stagnation among younger workers. A negative relationship exists between the rising wage share of older age cohorts and wage growth in younger groups (Ghilarducci et al., 2017). Older workers now hold a larger share of total earnings and employment, while younger cohorts experience slower wage advancement. These trends reflect a redistribution of wage growth towards older age groups during a period of demographic transition. Based on this, the research question of the study will be: “How does the increasing wage share of elderly workers in the Japanese labour force influence the wage growth of younger workers?”

Data source and overview

Dataset

The dataset used in this study is the “Zenkoku Shūgyō Jittai Panel Chōsa” (全国就業実態パネル調査), known as the National Employment Situation Panel Survey (JPSED) in English. The Survey is carried out annually by the Recruit Works Institute (リクルートワークス研究所) (Recruit Works Institute, 2025). The purpose of the dataset is to track employment conditions, income, and general living situations of people in Japan. The dataset gives a broad overview of the labour market and how it changes over time. The survey uses online monitoring methods and includes individuals aged 15 and above. To ensure representativeness, the sample is weighted to align with national distributions by gender, age group, region, job type, and educational background. The data includes details such as employment type (e.g., full-time, part-time, self-employed), income levels, working hours, and other basic lifestyle and demographic information. The Survey is also used to create the “Works Index,” which presents how people work through five different perspectives. This allows for a multi-dimensional understanding of labour trends in Japan. Overall, the JPSED offers a highly

detailed panel dataset, useful for analysing the structure and transformation of the Japanese workforce. It serves as an important source for research and policy development.

Model data

The data used covers years 2015 to 2024. The model data is aggregate panel data, with each row representing a year and an age group. Variables are expressed as attribute shares, which denotes how many people in that age group fall into a certain category of that specific independent variable. This format allows for an analysis of how group compositions changed over time across age cohorts (Recruit Works Institute, 2025).

Descriptive statistics

The dataset gives a detailed summary of income distribution across Japanese population groups and uses yearly wage brackets from 2015 to 2024. Each observation shows the percentage of people in a specific age group who fall into a certain income level. The descriptive statistics (mean, standard deviation, minimum, maximum, and median) show some important patterns about how wages are spread and how they have changed over time. Table 1 shows these basic descriptive results. Figure 1 shows how the share of age cohorts within selected wage brackets changed between 2015 and 2024. High-wage brackets (e.g., 10–12 million and 12 million+ Japanese Yen (JPY)) expanded in the older age groups (60–69, 70+), while younger groups (15–19, 20–29) shrank in low-wage brackets. This indicates an aging of the wage structure: older cohorts now receive a larger share of higher wages, while youth presence is shrinking in the lower brackets. Figure 2 compares the average population share in each wage bracket over time. The overall figure shows upward change: lower wage categories (under 2 million JPY) saw a decline in population share, while mid- and high-income categories increased slightly. This suggests an upward shift in wage concentration, with fewer people in low-income brackets and more in higher ones by 2024.

Table 1: *Descriptive statistics for salary, gender, and working hours variables*

Source: *Author’s own computation based on JPSED data.*

Variable	Description	Mean	SD	Min	Max	Median
Salary	< ¥500,000	0.1423	0.1331	0.0360	0.4800	0.0770
Salary	¥500,000– ¥1,000,000	0.1277	0.0589	0.0480	0.2670	0.1205
Salary	¥1,000,000– ¥2,000,000	0.1743	0.0465	0.1090	0.2770	0.1580
Salary	¥2,000,000– ¥3,000,000	0.1499	0.0348	0.0660	0.2430	0.1420
Salary	¥3,000,000– ¥4,000,000	0.1312	0.0526	0.0150	0.2050	0.1370
Salary	¥4,000,000– ¥5,000,000	0.0932	0.0503	0.0030	0.1770	0.0980
Salary	¥5,000,000– ¥6,000,000	0.0626	0.0418	0.0000	0.1380	0.0515
Salary	¥6,000,000– ¥7,000,000	0.0425	0.0313	0.0000	0.0850	0.0375

Salary	¥7,000,000– ¥8,000,000	0.0267	0.0235	0.0000	0.0670	0.0215
Salary	¥8,000,000– ¥9,000,000	0.0196	0.0182	0.0000	0.0580	0.0150
Salary	¥9,000,000– ¥10,000,000	0.0093	0.0092	0.0000	0.0320	0.0060
Salary	¥10,000,000– ¥12,000,000	0.0099	0.0085	0.0000	0.0300	0.0080
Salary	> ¥12,000,000	0.0107	0.0088	0.0000	0.0300	0.0090

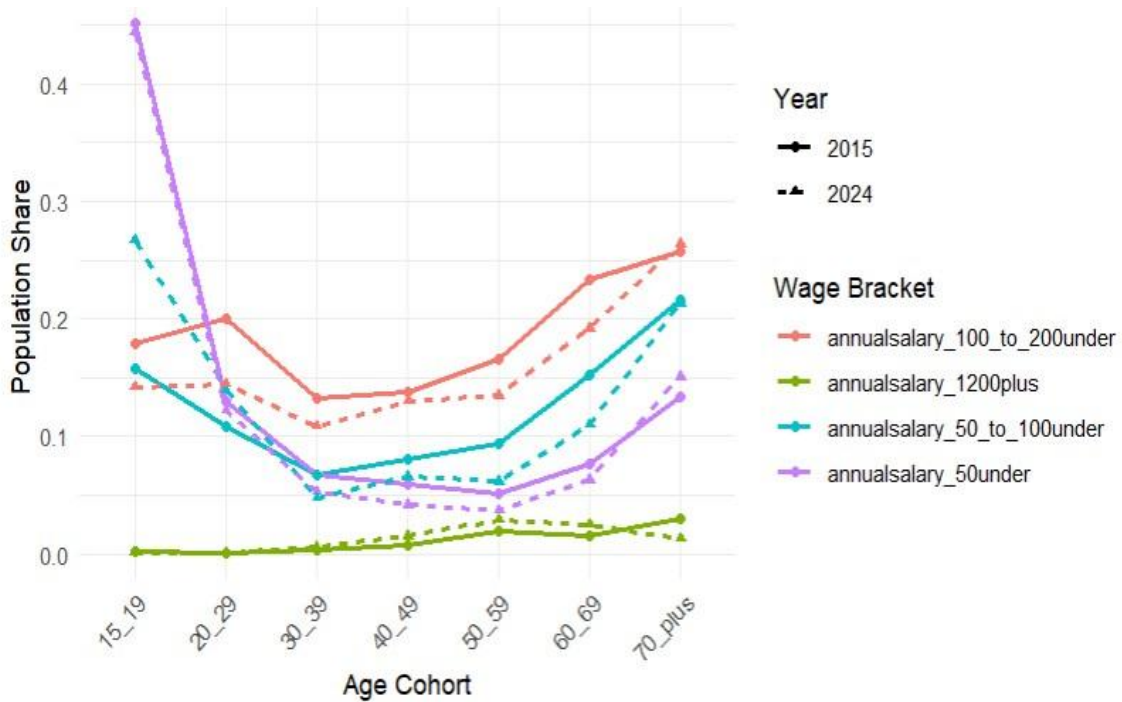


Figure 1: Change in wage bracket population shares by age cohort (2015 vs. 2024)

Overall Wage Score (dependent variable)

The Overall Wage Score is calculated by assigning numerical values to each wage bracket based on their midpoint values (e.g., 7.5 million JPY, 15 million JPY, etc.), and then computing a weighted average using the population share of each bracket.

The purpose of calculating the Overall Wage Score is to capture changes in the wage




Figure 2: Changes in the population share across annual wage brackets JPY, 2015–2024
Source: Author’s own computation based on JPSED data.

distribution of each age group over time. Instead of focusing on one specific bracket, the score summarizes the structure of the entire wage profile. It allows us to see whether an age group is gradually moving into higher or lower wage brackets and to track such movement consistently across years. The use of this score is appropriate for this specific research because it simplifies the wage composition into a continuous variable while still reflecting the underlying structure. It avoids issues that may arise from focusing on a single wage bracket, such as volatility or bracket shifts. This approach also provides a basis to test how changes in population composition by age and employment status influence the wage structure of the labour market over time. For observation i , the overall wage score is:

$$\text{Overall Wage Score}_i = \sum_{j=1}^J w_{ij} \cdot s_j \quad (1)$$

- i : Age group index (e.g., $i = \text{“30–39”}$)
- j : Wage bracket index
- J : Total number of wage brackets
- w_{ij} : Proportion of individuals in age group i falling into wage bracket j
- s_j : Score assigned to wage bracket j , representing midpoint or average wage
-

This reduces the full wage distribution to one numeric score. It keeps both wage level and spread and is useful for regression.

A more straightforward method would have been to use a full matrix structure, where each cell represents the wage score of a specific age group, wage bracket, and year. This setup could have made it easier to model how changes in population composition affect wage outcomes across the entire distribution. However, this approach was not chosen because of the limitations of the dataset. The data was already aggregated and given in percentage form. This does not include individual-level wage information or the tracking of the same people over time. Some age and wage bracket combinations appear with zero or missing values in certain years. Keeping the full matrix structure would require including all age groups and all wage brackets across all years. But many of these cells are sparse or unstable, especially in older and higher-income categories. To make the matrix usable, we would have to drop some age groups or brackets, which would reduce

the completeness of the analysis. Another issue is multicollinearity. The shares across wage brackets in a given year always total one. If one bracket increases, another must decrease. Including all of them in a model would create strong correlation between predictors, which leads to unstable estimates and inflated standard errors.

For these reasons, we use the Overall Wage Score instead. This variable compresses the full wage distribution into a single summary score for each age group and year. It allows us to study structural shifts in wage composition while avoiding the problems of missing data, overfitting, and multicollinearity.

In Japan, most firms impose a mandatory retirement age at 60, after which many workers are re-employed under short-term contracts. These new contracts typically offer lower wages. On average, companies using the continued-employment system report a wage drop of around 24% when workers turn 60. This institutional shift justifies separating the 50–59 age cohort from older groups, since workers above 60 often experience different employment terms and income levels (Moriyama, 2024). Figure 3 shows a clear and consistent gap between younger (15–19) and older (50–59) age groups in overall wage score. Over time, the 50–59 group steadily increases their score, while the 15–19 group stays low and flat. This suggests that older workers are moving into or maintaining high-income positions, while younger workers see no real improvement. The growing distance between the two lines over time supports the idea that the wage growth space for younger workers is narrowing.



Figure 3: *Development of overall wage score across years by age group*
Source: *Author’s own computation based on JPSED data.*

Hypothesis

Based on the literature review and descriptive statistics, we constructed the following hypothesis:

H1: As the population share of high-income cohorts increases relative to the 2015 baseline, the overall wage score decreases for younger age cohorts (e.g., ages 15–29).

Annual Salary (independent variables)

The independent variable, Annual Salary, as shown in Table 2, is subject to absolute and percentage changes in values across time, based on a fixed baseline year 2015. The logic behind the calculation can be expressed in the following way.

Assume that $x_{j,t}$ indicates the value of a variable for a specific age group j in year t . The corresponding baseline value for age group j in the year 2015 is denoted as $x_{j,2015}$. Then, the absolute change is calculated as:

$$\Delta x_{j,t} = x_{j,t} - x_{j,2015} \quad (2)$$

This result shows the direct increase or decrease in the variable from its 2015 level. It measures the magnitude of difference without standardizing by the initial value.

In addition to this, the percentage change is computed. This is defined as:

$$\% \Delta x_{j,t} = \frac{x_{j,t} - x_{j,2015}}{x_{j,2015} + \varepsilon} \times 100 \quad (3)$$

Here ε is a very small constant (usually set to 10^{-6}), which is added to avoid division by zero if $x_{j,2015}$ happens to be zero. This formula expresses the change in relative terms, and indicates the proportional difference as a percentage compared to the baseline. Age is factorized in the model to reflect group-level differences not captured by a linear numeric form. Wage patterns do not increase linearly with age: young, middle-aged, and older workers are in different labour market positions. By treating age as a factor, the model estimates a separate effect for each group without assuming equal distance between them. This allows better control for life-cycle effects, such as early-career low wages, mid-career peaks, and possible late-career declines. Since the dependent variable is already age-specific (wage score per age group), methodologically it also makes sense to use age as a categorical variable. This helps capture structural differences across cohorts without forcing a single slope on age.

Even though Annual Salary was in percentage form, standardization was still needed. This is because the spread of different attribute values was very different across variables.

Table 2: Variable attributes and types
Source: Author's own elaboration based on JPSED data.

Variable	Attributes (English)	Value Type
<i>Age</i>	15–19, 20–29, 30–39, ..., 70+	Categorical
<i>Year</i>	2015 to 2024	Integer
<i>Annual Salary</i>	Under 0.5 million JPY	Percentage
	0.5–1 million JPY	Percentage
	1–2 million JPY	Percentage
	2–3 million JPY	Percentage
	3–4 million JPY	Percentage
	4–5 million JPY	Percentage
	5–6 million JPY	Percentage
	6–7 million JPY	Percentage
	7–8 million JPY	Percentage

	8–9 million JPY	Percentage
	9–10 million JPY	Percentage
	10–12 million JPY	Percentage
	12 million JPY or more	Percentage
	Respondents Count	Count

For example, the standard deviation of 0.5–1,0 million JPY was around 0.133, while 10–12 million JPY had a much smaller spread, around 0.0085. This means that the same percentage unit meant very different variability. Variables with wider spread could dominate the model just because of their larger scale, and not because they are more important predictors. This also helps with model stability, especially when there are interaction terms or multicollinearity. Furthermore, mixed models are sensitive to differences in scale. Without standardization, variables with large spreads may cause estimation problems or convergence issues. Standardization helps avoid these. Also, standardized coefficients are easier to compare. Since all predictors are on the same scale, coefficient size reflects relative effect more clearly. So, even with only percent units, standardization improves balance, comparability, and model stability.

Baseline model

Modelling strategy justification

Linear mixed modelling was used at the beginning. In theory, linear mixed-effects models are more suitable than the Ordinary Least Squares (OLS) method, when data has a grouped structure (like ours, where observations are nested within years). OLS treats all observations as independent, but that ignores the fact that wage outcomes in the same year may be more similar than across years. `lmer()` includes random intercepts for Year, which accounts for unobserved heterogeneity across time. This is useful when it is expected that yearly changes in policy, economy, or population might shift overall wage levels. Mixed models also help reduce bias in fixed effect estimates by controlling for these group-level effects. However, when bias was estimated using Linear mixed-effects modelling functions in R (`lmer()` from `lme4` package), with the structure `Overall_Wage_Score ~ ... + (1|Year)`, it gave a singular fit. That means the variance of the random effect for Year was estimated as zero or very close to zero. In such cases, the model cannot separate between-group and within-group variation. This thus suggests that the grouping structure is not supported well by the data. The Robust version of Linear mixed-effects modelling from R studio (`rlmer()`) is also used and there was several changes—removing variables, changing fixed effects, adding polynomial terms for Year, and even grouping years—but none of these solutions fixed the singularity problem. The model remained unstable, and the results could not be trusted. The frequentist tools dropped the random effect silently or gave convergence warnings. Because of this, it was then changed to Bayesian modelling using the Bayesian Regression Models using Stan `brms` R package. In Bayesian estimation, prior distributions help the model estimate random effects even when group-level variation is small. Instead of collapsing, the model keeps the random effect and gives a full posterior distribution. This allowed me to include Year as a grouping factor without singularity problems. The Bayesian model also gave better diagnostics and interpretability. In the end, the choice to use Bayes was not just preference but a necessary step. The data structure did

not support random effects in lmer or rlmr, but the Bayesian method allowed me to keep the model structure and to estimate the effects more reliably.

In this study, only the salary-related variables were used in the final model. This decision was based on the complexity of the model. Although the JPSED dataset includes many other variables like working hours, contract type, or job satisfaction, the inclusion of all of these in the model made the model unstable and hard to interpret. The model was already difficult to fit, and it needed to stay focused and reliable. The Annual Salary also provided good detail. There are 13 income brackets, covering the full range from very low to very high wages. Using these brackets, a weighted wage score for each group was created, which kept both the level and the spread of income. This gave a strong and clear outcome variable for the model. The main research question is about wage pressure and age structure. So it made sense to focus directly on salary changes. Adding other variables would have made the model harder to estimate and less clear in terms of interpretation. Because a Bayesian model was used, it was also important to reduce complexity and avoid variables that could cause convergence problems or unstable posteriors.

Model selection procedure

Likelihood nesting was used to compare different model forms and to find the best structure for estimating overall wage score. I started out from a basic model with just the intercept, with additional variables one by one. These included changes in population shares across salary brackets, contract types, industries, and gender composition. Each new model was compared to the previous using Bayesian metrics like expected log predictive density (ELPD) and leave-one-out cross-validation (LOO). Among the tested models, the one with only the standardized change in the population share of those earning between 10 and 12 million yen, along with a random intercept for *Year*, showed best performance. This model showed better predictive fit and convergence than more complex models, and did not include unnecessary parameters.

Final model setup

In the analysis of Wage Scores, it is important to consider that wages change over time. Each year may have different average wages because of economic conditions or policies. To capture this, the model lets the baseline wage change for each year. The following Bayesian multilevel model was estimated to evaluate the impact of changes in high-income population share on overall wage outcomes:

$$y_{i,t} \sim N(\mu_{i,t}, \sigma^2) \quad (4)$$

$$\mu_{i,t} = \beta_0 + \beta_1 \cdot \Delta x_{j,t} + b_{0[t]} \quad (5)$$

$$b_{0[t]} \sim N(0, \tau^2) \quad (6)$$

- $y_{i,t}$ is the observed *Overall Wage Score* for age group i in year t ,
- $\Delta x_{j,t}$ is the standardized change in population share for wage bracket j (e.g., those earning 10–12 million yen) in year t relative to 2015,
- $\mu_{i,t}$ is the expected wage score given the covariate and year,
- β_0 is the global intercept (grand mean of the wage score),
- β_1 is the slope coefficient capturing the effect of wage bracket population change,
- $b_{0[t]}$ is the random intercept for year t , modelled as normal deviation with variance τ^2 ,
- σ^2 is the residual variance not explained by the model.

- The Bayesian estimation provides posterior distributions for the parameters:

$$p(\beta_0, \beta_1, \sigma, \tau | y) \propto p(y | \beta_0, \beta_1, \sigma, \tau) \cdot p(\beta_0, \beta_1, \sigma, \tau)$$

Priors are set by default as weakly informative within the brms package to ensure numerical stability while allowing the data to drive inference. Using Bayesian methods lets us add prior information about parameters. This is useful when data are limited or noisy. Priors help make estimates more stable and less prone to overfitting. Bayesian inference uses Markov Chain Monte Carlo (MCMC) sampling methods like No-U-Turn Sampler (NUTS) to explore complex posterior shapes. This ensures that the final results for wage scores and yearly effects are reliable and supported by the data.

Research results

Hypothesis test

The model examines how changes in the population share of high-income earners—specifically those earning between 10 to 12 million yen—relate to overall wage score using a Bayesian multilevel framework. This approach includes random intercepts for year to account for variation across time. The intercept is estimated at 299.18, which represents the expected wage score when the independent variable is at its average. This serves as a reference point for interpreting the effect of salary bracket changes. The main independent variable, the standardized change in the population share of people earning 10–12 million yen, has a coefficient of -41.97 . This value is credibly below zero, with a 95% credible interval ranging from -68.39 to -16.32 . This suggests that as the proportion of high-income earners increases, the wage score tends to decrease. This supports the hypothesis that expansion in the upper salary bracket may contribute to wage compression or reduce wage space available to other groups. The model also includes a random intercept for year. The standard deviation of the intercepts across years is 14.23, with the credible interval $[0.55, 44.81]$, which indicates that wage levels vary moderately by year, possibly due to macroeconomic shifts, inflation or labour policy changes. The residual standard deviation is estimated at 104.77, which shows that most variation in wage scores is not explained by the model. Still, the model captures meaningful structure.

Table 3: *Bayesian regression results for overall wage score*
Source: *Author's own elaboration based on JPSED data.*

Parameter	Estimate	Est. Error	95% CI	Rhat
β_0 (Global Intercept)	299.18	14.37	[270.58, 327.36]	1.00
$\beta_1 \times \Delta x_{j,t}$ (Pop. change in 10–12 million JPY bracket)	-41.97	13.10	$[-68.39, -16.32]$	-1.00
$SD(b_{0[Year]})$ (Variation in intercept across years)	14.23	12.05	[0.55, 44.81]	1.00
σ (Unexplained Variance)	104.77	9.81	[87.82, 126.39]	1.00

The results, which are summarised in Table 3, supports the hypothesis that increasing high-income concentration distorts wage structure. Although these workers earn more individually, their

rising population share coincides with a decline in the composite wage score, which implies structural wage rigidity or displacement effects elsewhere. This suggests that retaining high-wage, older employees may prevent or hinder wage growth for other groups—particularly younger or mid-income workers—, which creates an imbalance in labour market opportunities over time.

Convergence checks

As shown in Figure 4, the traceplots for all parameters indicate good mixing across the four sampling chains. The chains overlap consistently and move densely throughout the posterior distribution without getting stuck or drifting apart. Each chain appears to stabilize early and remains centred around the same region of the parameter space, which suggests that the sampler has adequately explored the distribution. No strong autocorrelation or irregular patterns are observed, and the overall appearance is typical of well-behaved MCMC output. These results support the conclusion that the model has converged properly, and the posterior draws are reliable for interpretation. With all diagnostic values falling within acceptable ranges, the potential scale reduction factor (Rhat) is equal to 1.00 for every parameter, including both fixed effects and variance components, which suggests that the Markov chains converged and mixed well. There were no divergent transitions during the sampling process, which means the algorithm moved smoothly through the posterior space without hitting regions that would cause instability. In other words, the sampler did not get stuck or behave erratically while drawing samples. Also, none of the iterations reached the maximum allowed tree depth. This is important because when the sampler uses its full path length, it suggests the posterior may be too complex or that the model is poorly specified. Since that did not happen, it can be said that the sampling paths were efficient and the geometry of the model was manageable. The energy diagnostic, which tests whether the sampler explores the space evenly, showed no issues. If this test had failed, it would have meant that the sampler was inefficient or imbalanced, but here the energy levels behaved as expected. Together, these results mean the sampling process was stable, efficient, and well-suited to the structure of the model. The posterior estimates can therefore be treated as reliable.

Posterior predictive fit

The posterior predictive check shown in Figure 5 gives a visual means to assess how well the model can reproduce the distribution of the observed outcome. The dark blue line shows the density of the actual observed data (y), while the lighter lines represent simulated draws (y_{rep}) from the posterior predictive distribution.

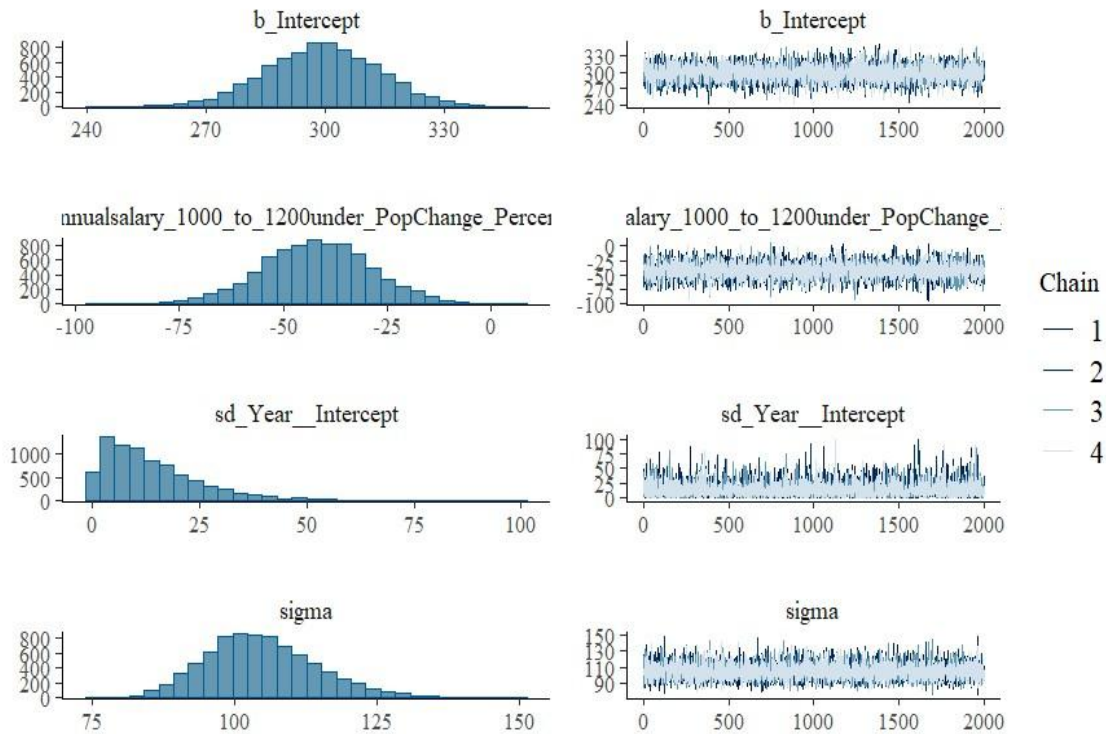


Figure 4: *Posterior histogram / Marginal distribution and traceplot*
Source: *Author's own results.*

In general, the model seems to capture the main shape and spread of the data including the clear bimodal structure. This suggests that the model does a reasonable job at reflecting the overall data-generating process. Still, the simulated densities appear a bit more peaked in the centre and smoother in the tails when compared to the observed line. This shows that the model might be slightly underfitting the sharper details in the observed distribution. One possible reason for this is that the model assumes normally distributed residuals, which might not fully match the shape of the actual wage score distribution, especially if the real data has heavier tails or more sudden changes in density. This can cause the replicated outcomes to be too smooth or to miss out on sharper peaks or troughs as shown by the observed data.

Therefore, while the model's predictive fit is mostly acceptable and follows the general trend of the data, the small mismatch in peaks and tails suggests that the model might be missing some complexity. The model may improve by testing other error distributions or adding predictors that can explain more of the variation across groups or values. This can be visually confirmed on Figure 5, which shows that the observed density (y) and the posterior predictive draws (y_{rep}) are close in general form but differ in finer details like peaks and tails.




Figure 5: *Posterior predictive distribution plot*
Source: *Author's own results.*

The posterior distribution shown in Figure 6 corresponds to the estimated coefficient for the standardized change in the population share earning between 10 and 12 million yen annually. The density plot summarizes the uncertainty around this effect based on the posterior draws from the Bayesian model. The distribution is uni-modal and is centred around approximately -43 , with most of its mass falling between -68 and -16 . The shaded region represents the 95% credible interval, which lies entirely below zero. This indicates that the model assigns high probability to the coefficient being negative, and thus there is strong Bayesian evidence that an increase in the share of individuals in this high-income group is associated with a decrease in the overall wage score. The posterior is fairly symmetric and tightly concentrated, which implies that the model estimates this effect with relatively high precision. Because the interval does not cross zero and the density is clearly distinct from it, the effect can be considered statistically credible in the Bayesian framework. This finding supports the hypothesis that growth in the high-income cohort may contribute to wage score decline, possibly due to structural wage imbalance or shifts in population composition. We can see this more clearly in Figure 6, which visualizes the marginal posterior distribution for the 10–12 million JPY range and highlights the model's high confidence in a negative effect.




Figure 6: *Marginal posterior distribution for 10–12 million JPY*
Source: *Author’s own results.*

Overfitting or overshrinkage

The random intercept for Year has an estimated standard deviation of approximately 14.23, with a 95% credible interval ranging from about 0.55 to 44.81. This suggests that there is some variation in the overall wage score across different years, and while the estimate is relatively uncertain (as shown by the wide interval), the lower bound is well above zero. This means that year-to-year differences are likely not negligible, and the model accounts for them by including Year as a group-level effect. The fact that the interval does not include zero indicates that the model treats these differences as meaningful data rather than as noise. The residual standard deviation, which reflects the amount of variability in the wage score not explained by the predictors or the grouping structure, is estimated at around 104.77. Its credible interval is relatively narrower, ranging from approximately 87.82 to 126.39, which indicates more certainty in this estimate. This value represents the typical deviation of individual observations from their group-specific predicted values. In this case, it shows that while the model captures some structure across years, a large portion of the variance remains unexplained at the individual level. In sum, the model recognizes that some variation in wage scores is due to differences across years, but the majority of the variation still lies within years and is not explained by the predictors included. The hierarchical structure captures both components, but the residual error remains dominant, which suggests potential room for improving the model’s explanatory power with additional predictors or refinements.

Discussion

This study used publicly available aggregate panel data, which limited the range of variables that could be included in the model. Because of the structure of the dataset, important individual-level

factors such as work hours, job position, and tenure could not be used. These elements are highly relevant in explaining wage dynamics, especially in the Japanese employment system, where roles and internal career progression matter. As a result, the model focuses only on the effect of population change within specific wage brackets and does not account for these additional drivers. From 2015 to 2024, descriptive trends indicate that older age groups (60–69 and 70+) increasingly occupy high-income brackets, especially in the 10–12 million yen and 12 million yen-plus ranges. Meanwhile, younger cohorts are shrinking in lower-income categories. This reflects an aging wage distribution, where wage growth is concentrated among older workers, possibly narrowing future mobility for younger employees. The main result supports the hypothesis. The standardized change in the population share of the high-income elderly (10–12 million yen) has a negative, statistically credible coefficient of -41.97 . This implies that as more older individuals move into upper wage tiers, the overall wage score declines. The result suggests that this demographic shift compresses wages and may crowd out other groups, especially younger workers. The Bayesian multilevel model used here offers stable estimates, which overcome instability issues present in traditional mixed models with limited variation across groups.

Although the model showed clear results regarding the association between an increase in high-income earners and overall wage outcomes, the variance across years was not significant. This implies that other unexplained differences may stem from individual-level variations that aggregate data simply cannot capture. In the future, as we understand these effects better, we can use the model to produce a leading indicator. This would help predict how wage trends might shift—especially if wage growth for younger groups is being suppressed—and act early to prevent imbalance.

Future research will incorporate individual-level panel data from the Japanese Panel Study of Employment Dynamics (JPSED), which enables the development of a more complex multilevel model that includes additional variables. Such a model would allow for a more detailed understanding of the mechanisms behind wage shifts in an aging society. The topic remains of intense interest and will continue to be explored using improved data in subsequent stages.

References

Abegglen, James C. (1958). *The Japanese Factory: Aspects of Its Social Organization*. *The Journal of Asian Studies*, 18(1), 138–139. <https://doi.org/10.2307/2941308>

Ghilarducci, T., Papadopoulos, M., & Radpour, S. (2017). *Relative Wages in Aging America: The Baby Boomer Effect*. Schwartz Center for Economic Policy Analysis and Department of Economics, The New School for Social Research, Working Paper Series 2017-5.

MHLW (Ministry of Health, Labour and Welfare). (2024). 令和 6 年版厚生労働白書（令和 5 年度厚生労働行政年次報告）. Retrieved from <https://www.mhlw.go.jp/stf/wp/hakusyo/kousei/24/index.html>. Accessed 23 June 2025.

Mohnen, P., & Strauss, A. (2024). *The Impact of the Retirement Slowdown on the U.S. Youth Labor Market*. Retrieved from <https://doi.org/10.1086/725874>

Moriyama, Tomohiko. (2024). *The Retirement Process of Older Workers in Japan: Perspectives from Legislative Frameworks and Human Resource Management*. The Japan Institute for Labour Policy and Training.

https://www.jil.go.jp/english/publications/documents/jpt/2024/JPT_074.pdf. Accessed 27 June 2025. Recruit Works Institute. (2025). 全国就業実態パネル調査 [*National Employment Reality Panel Survey*]. Works-i. <https://doi.org/10.34500/SSJDA.1730>

日本労働組合総連合会. (2023). 連合・賃金レポート 2023. Retrieved from https://www.jtuc-rengo.or.jp/activity/roudou/shuntou/2024/wage_report/wage_report.pdf. Accessed 23 June 2025.