RIETI Discussion Paper Series 25-E-081

Nonlinear Earnings Dynamics and Inequality over the Life Cycle: Evidence from Japanese Municipal Tax Records

KITAO, Sagiri
RIETI

SUZUKI, Michio
Tohoku University

YAMADA, Tomoaki
Meiji University

The Research Institute of Economy, Trade and Industry
https://www.rieti.go.jp/en/

RIETI Discussion Paper Series 25-E-081

August 2025

Nonlinear Earnings Dynamics and Inequality over the Life Cycle: Evidence from Japanese Municipal Tax Records*

Sagiri KITAO, Research Institute of Economy, Trade and Industry

Michio SUZUKI, Tohoku University

Tomoaki YAMADA, Meiji University

Abstract

This paper examines life-cycle earnings risk and income inequality in Japan using municipal administrative tax records covering the period 2011–2021. We estimate an age-dependent quantile model that decomposes idiosyncratic earnings into persistent and transitory components, allowing persistence to vary nonlinearly with individuals' positions in the earnings distribution and the size of shocks. We find that persistence is high for shocks consistent with an individual's earnings history but falls sharply for "reversal" shocks, which may represent career changes. Cross-sectional analysis shows that households pool income effectively: equivalized household income displays lower dispersion and a J-shaped life-cycle inequality profile compared to a monotonically rising profile for individual earnings. However, impulse response analysis reveals that when individuals or households at a given percentile of the persistent component distribution receive either a high or low percentile draw from the innovation distribution, the resulting earnings changes are larger for equivalized household earnings than for the earnings of the household head alone. This indicates that household and individual earnings distributions have distinct dynamic properties, with household-level responses potentially reflecting correlated spousal shocks, joint labor supply decisions, and demographic adjustments.

Keywords: earnings dynamics, income inequality, life-cycle, administrative data, nonlinear panel quantile regression, intra-household risk sharing

JEL classification: J31, D31, C23

* This study is conducted as a part of the Project "Household Heterogeneity: individuals, families and macroeconomy" undertaken at the Research Institute of Economy, Trade and Industry (RIETI). We are grateful for helpful comments from participants of the international conference "Economic Research using Administrative Data". We also thank Kyoji Fukao, Keiichi Kobayashi, and Eiichi Tomiura for their valuable comments at a seminar hosted by RIETI. We acknowledge support from JSPS KAKENHI grant J22H05008.

1 Introduction

Understanding idiosyncratic earnings risk is fundamental to macroeconomics, labor economics, and public finance. While large administrative panel datasets have been extensively utilized internationally to analyze this risk, comprehensive administrative data in Japan has remained largely untapped.[1] Therefore, exploiting local administrative data is a crucial step toward gaining a deeper and more granular understanding of idiosyncratic earnings risk in Japan, complementing existing evidence from survey data.

This paper takes this step by leveraging newly anonymized local tax micro-data provided through the Local Government Data Development Project, a partnership between participating municipalities and the Center for Research and Education in Program Evaluation (CREPE) at the University of Tokyo. Our primary contributions are threefold. First, we document the evolution of cross-sectional dispersion measures for labor earnings, total income, and after-tax income, both over time and over the life cycle. We find that dispersion measures for equivalized household earnings and income consistently lie below those for the household head's earnings, suggesting effective income smoothing at the household level. Second, we apply the flexible quantile-based model of earnings dynamics developed by Arellano et al. (2017) (hereafter, ABB) to these Japanese data. Our findings confirm the presence of significant nonlinear persistence; for instance, the predictive power of past earnings drops sharply for individuals experiencing "microeconomic disasters," such as a high-earning individual receiving a large negative shock. Finally, building on this model, our impulse response analysis reveals that equivalized household earnings are more responsive to persistent shocks than are the household head's earnings alone.

The data for this study are drawn from the administrative local tax records held by a municipality in Japan. While not representative of Japan as a whole, the dataset's administrative origin ensures comprehensive coverage of the local resident population and virtually eliminates the sampling error, attrition, and under-reporting issues that often affect survey data. Crucially, the data contain unique individual and household identifiers, allowing us to construct a large panel that makes the estimation of a rich model feasible. This structure enables us not only to track individual earnings trajectories but also to aggregate income at the household level to examine intra-household insurance mechanisms. Furthermore, we construct after-tax income measures by accounting for local and national taxes, which allows for a direct assessment of the Japanese tax system's role in mitigating income volatility. A limitation of our data, however, is that it does not contain information on public transfers.

[1] See Guvenen et al. (2022) for administrative data made available for research use in 13 countries included in the Global Repository of Income Dynamics (GRID).

Recently, the Japanese National Tax Agency initiated a program to provide national income tax data for statistical research purposes. However, the coverage of this national tax data is inherently limited because a significant fraction of Japanese employment income earners pay taxes exclusively through withholding tax, meaning their earnings are not captured in the final tax return data.

Methodologically, we adopt a flexible, quantile-based panel approach to model earnings dynamics, building directly on the nonlinear framework developed by ABB. We first decompose the idiosyncratic component of log earnings—defined as the residual from a regression on age and year dummies while controlling for individual fixed effects—into a persistent and a transitory component. Following ABB, we model the conditional distributions of these components flexibly using quantile functions that depend on lagged earnings levels and age. The model is estimated nonparametrically by fitting these quantile functions at various percentiles on a pre-specified grid, with linear interpolation used between grid points.

A key advantage of this approach is its ability to generate rich, state-dependent dynamics. We examine conditional moments derived from the estimated quantile functions. In particular, persistence, defined as the partial derivative of the conditional quantile function with respect to past earnings, is not a single parameter as in canonical linear autoregressive models. Instead, our flexible specification allows persistence to vary with an individuals past earnings level, the specific percentile of the shock, and age.

Our empirical analysis first documents the evolution of key inequality measures such as the variance of logarithms, the Gini coefficient, and percentile ratios (P90/P50, P50/P10) over time and across the life cycle. We find that these dispersion measures are consistently lower for equivalized household earnings than for household head earnings, suggesting a degree of income smoothing at the household level. Furthermore, after-tax income measures exhibit lower inequality than their before-tax counterparts, highlighting the redistributive role of the tax system.

Turning to earnings dynamics, our estimation of the ABB model reveals significant nonlinear persistence. The predictive power of past earnings remains high for individuals experiencing shocks consistent with their recent history, but it drops sharply for those hit by large 'reversal' shocks, such as a high-earner experiencing a major negative shock. This pattern points to the importance of events like career transitions or significant job displacements.

To quantify the impact of these dynamics, we compute impulse responses by simulating earnings paths starting at age 35. For individual-level analysis, we track individuals at different initial percentiles (10th or 90th) of the persistent component distribution. For household-level analysis, we track households whose heads are aged 35 and at different initial percentiles. At age 36, we apply either a 10th or 90th percentile draw from the innovation distribution, compared to a baseline 50th percentile draw, and track the resulting earnings differences over the life cycle. The analysis shows that these percentile transitions in the innovation distribution generate larger earnings changes for equivalized household income than for the household head's earnings alone. This difference in impulse responses reveals that household and individual earnings distributions have distinct dynamic properties, with the same percentile innovation having greater economic impact at the household level. This pattern likely reflects several mechanisms: correlated shocks affecting multiple household members (such as industry-specific downturns impacting both spouses), joint labor supply decisions in response to adverse events (such as caregiving needs), and demographic changes that alter household composition and the equivalization factor. However, it is important to note that our data lack information on public transfers, which may lead us to overstate the magnitude of earnings changes at the household level. Events that reduce household earnings-such as childbirth, illness, or eldercare responsibilities-are often partially offset by public support programs, and the absence of these transfers in our analysis may exaggerate the apparent volatility of household income dynamics.

Our work is most closely related to ABB, whose nonlinear panel data framework we adopt. While they used PSID survey data and Norwegian administrative data to jointly model earnings and consumption, our focus is on applying their earnings model to the unique context of Japanese administrative data. The large sample size allows us to examine the age-profile of conditional moments in greater detail. Our analysis also connects to recent extensions of this framework, such as Arellano et al. (2024), which incorporated unobserved household heterogeneity and advanced computational methods to study consumption. De Nardi et al. (2020) has embedded the ABB earnings process into a large-scale life-cycle model to quantify the welfare impact of nonlinear risk. While a full welfare analysis is beyond the scope of this paper, we provide suggestive evidence by calculating the cumulative effects of persistent shocks based on our impulse response estimates.

Our analysis of household-level dynamics relates to studies such as Blundell et al. (2015), which also uses administrative data to disentangle insurance mechanisms. Our work differs on several key dimensions. While their Norwegian data contains comprehensive information on both taxes and public transfers, our Japanese data provides detail on local inhabitant tax but not transfers. Methodologically, we adopt the flexible, nonlinear framework of ABB, which allows for state-dependent persistence. In contrast, Blundell et al. (2015) extend a canonical linear model by incorporating heterogeneous income profiles and allowing the variances of shocks to vary with both age and time, with a primary focus on documenting how these dynamics differ across education groups. Blundell et al. (2015) find that spousal income provides little additional insurance on top of the Norwegian tax and transfer system, whereas our household-level analysis reveals that the same percentile changes lead to larger responses in equivalized household earnings than in household-head earnings.

The remainder of this article is organized as follows. Section 2 describes our model of idiosyncratic earnings dynamics. Section 3 explains our data. Section 4 documents the evolution of cross-sectional dispersion measures of various income variables over time and over life-cycle. Section 5 presents quantile-based conditional moments of idiosyncratic components of the income variables. Section 6 reports the estimation results of the ABB model and the associated impulse response analyses. Section 7 concludes. Appendix A provides a more detailed description of the data used in this paper. Appendix B provides additional results including robustness checks using data from an alternative municipality.

2 A Model of Idiosyncratic Earnings Dynamics

Let Y_{it} denote before-tax labor earnings of individual i at time t. We model $\ln Y_{it}$ as a function of a full set of age dummies and individual and year fixed effects. Thus,

$$\ln Y_{it} = \sum_{k=2}^{A} \beta_k I(age_{it} = k) + \sum_{s=3}^{T} \gamma_s I(year_t = s) + \mu_i + y_{it}, \tag{1}$$

where y_{it} represents the residual. We consider y_{it} as the variable idiosyncratic component of individual earnings and thus examine its distribution in detail in this paper.

Our specification in equation (1) deviates from the standard approach in the earnings dynamics literature in two important ways. First, we include individual fixed effects μ_i rather than observable demographic characteristics as controls. The standard practice, also employed by ABB, is to include controls for education, race, region of residence, family composition, and other demographic characteristics. However, our administrative tax records do not contain information on educational attainment or other key demographic variables that are typically available in survey data. Given this data imitation, we opt to include individual fixed effects, which capture all time-invariant unobserved heterogeneity, including education, ability, cohort effects, and other permanent individual characteristics that affect earnings levels.

Second, the inclusion of individual fixed effects μ_i along with age and year dummies creates a multicollinearity issue, since cohort effects are now absorbed by the individual fixed effects (as cohort = year of birth = year − age for each individual). To resolve this identification problem, we drop the year dummies for $t = 1, 2$ rather than imposing the linear constraints typically used in the literature. This normalization allows us to identify the remaining parameters while maintaining the comprehensive control for individual heterogeneity provided by the fixed effects.

The subsequent decomposition of the idiosyncratic component y_{it} into a persistent component η_{it} and a transitory component ϵ_{it} follows the nonlinear panel data framework developed by ABB. This approach models the evolution of the persistent component η_{it} as a general Markov process, defined by the conditional quantile function $Q_\eta(\eta_{it-1}, age_{it}, u_{it})$, where u_{it} is a uniformly distributed innovation. Similarly, the transitory shock ϵ_{it} and the initial condition η_{i1} are modeled via age-dependent quantile functions Q_ϵ and Q_{η_1}

respectively.

$$
\begin{cases}
y_{it} = \eta_{it} + \varepsilon_{it}, \\
\eta_{it} = Q_\eta(\eta_{it-1}, age_{it}, u_{it}), \quad (u_{it}|\eta_{it-1}, \eta_{it-2}, \ldots) \sim \text{Uniform}(0,1), \\
\varepsilon_{it} \sim Q_\varepsilon(age_{it}, u_{it}^\varepsilon), \quad u_{it}^\varepsilon \sim \text{Uniform}(0,1), \\
\eta_{i1} \sim Q_{\eta_1}(age_{it}, u_{it}^{\eta_1}), \quad u_{it}^{\eta_1} \sim \text{Uniform}(0,1).
\end{cases}
\tag{2}
$$

Q_η, Q_ε, Q_{η_1} represent quantile functions. A key feature of this model, adopted from ABB, is its ability to capture nonlinear dynamics. Unlike standard linear models where the persistence of shocks is constant, this framework allows the persistence of past earnings components (η_{it-1}) to vary depending on the level of past earnings, the size and sign of the current innovation (u_{it}), and other covariates such as age_{it} via the shape of the quantile function Q. This nonlinearity allows the model to accommodate asymmetric effects and situations where large shocks (e.g., associated with job loss, health events, or career changes) might alter the influence of past earnings history. Capturing the full conditional distribution of the persistent component is crucial for understanding earnings risk and its implications.

As ABB show, the nonparametric identification of the earnings model in equation (2) relies on results from the measurement error literature, particularly Hu and Schennach (2008) and Wilhelm (2015). These papers establish that nonlinear models with latent variables can be identified under conditional independence restrictions, which naturally arise in our panel setting where earnings components are independent over time conditional on the persistent state. With annual data and at least four consecutive periods, the model's components—the Markov transition Q_η, transitory distribution Q_ε, and initial condition Q_{η_1}—are nonparametrically identified.[2] We impose time-stationarity, assuming that $Q_\eta(\cdot,\cdot)$ and the distribution of ε_{it} do not vary over calendar time. This allows us to pool observations of the same age across different cohorts, providing sufficient data for nonparametric identification at each age.

Following ABB and Arellano and Bonhomme (2016), we estimate the model using an iterative simulation-based algorithm that exploits the computational simplicity of quantile regression. The algorithm alternates between two steps:

Step 1 (E-step): Given current parameter estimates, draw values of the latent persistent components η_{it} from their posterior distribution conditional on observed earnings data, using a Metropolis-Hastings sampler.

Step 2 (M-step): Update the quantile function parameters by running standard quantile regressions: (i) regress η_{it} on $\eta_{i,t-1}$ and age to update Q_η; (ii) regress $y_{it} - \eta_{it}$ on age to update Q_ε; and (iii) regress η_{i1} on age to update Q_{η_1}. We estimate the

[2] See ABB Section 4.1 and Wilhelm (2015) for technical details.

coefficients for each quantile independently and apply a post-processing rearrangement following Chernozhukov et al. (2010) to avoid quantile crossing.

In practice, to avoid local optima, we first perform a parameter initialization step by estimating the model on multiple random subsets of the data (20% of the full sample each). We then use likelihood-weighted averages of these subset estimates as starting values for the full-sample estimation. We iterate the algorithm 200 times, with 200 Metropolis-Hastings draws in each iteration. We monitor convergence through the log-likelihood and compute final estimates as averages over the latter half of the iterations to account for simulation uncertainty. The quantile functions are approximated using piecewise-linear splines on a grid of quantiles, where at each quantile grid point, the quantile function is approximated by Hermite polynomials. Specifically, for Q_η, we use products of Hermite polynomials in $\eta_{i,t-1}$ and age, while for Q_ε and Q_{η_1}, we use Hermite polynomials in age. Beyond the quantile grid, we specify exponential (Laplace) tails to ensure proper behavior in the extremes. This approach ensures computational tractability while maintaining flexibility to capture nonlinear dynamics.

3 Data Description

Our empirical analysis draws on anonymized, individual-level panel data from administrative local tax records in a municipality in Japan, covering the years 2011–2021. The data were obtained through the Local Government Data Development Project, a collaboration between participating municipalities and the Center for Research and Education in Program Evaluation (CREPE) at the University of Tokyo.

The dataset includes detailed information necessary to compute local inhabitant tax liabilities, including income sources, statutory deductions, and basic demographic attributes such as age and gender. A key strength of the data is its administrative origin, which ensures full coverage of the resident tax population, including both individuals whose income is subject to the withholding tax system and those who file a final tax return. Each record contains a unique individual identifier and municipality code, enabling the construction of a multi-year panel.

A distinctive feature of the dataset, rare among administrative tax records, is the inclusion of a household identifier and each individuals relationship to the household head. This structure enables the construction of a household-level panel and allows us to identify household heads, spouses, and dependents and to analyze intra-household dynamics such as income aggregation and tax incidence.

We examine three main income measures: (1) labor earnings (pre-tax employment income), (2) total (net) income (sum of all taxable income sources subject to comprehensive taxation, excluding separately taxed interest and dividend income), and (3) after-tax income. All monetary values are deflated by the Core CPI (excluding fresh food, 2020 base year). For the household-level analysis, household earnings and income are equivalized by dividing by the square root of household size to account for economies of scale in consumption. The calculation methodology for after-tax income is detailed in Appendix A. To isolate the idiosyncratic component of each income measure, we regress the logarithm of income on a full set of age and year dummies, along with individual fixed effects (μ_i), as specified in equation (1).

For our analysis of idiosyncratic earnings risk, we aim to select a sample of individuals and households with stable attachment to the labor market, while retaining major changes in earnings and income, such as job loss. To achieve this goal, we follow Guvenen et al. (2021) and conduct sample selection as follows. For our individual-level analysis, we select individuals who meet the following criteria: 1. residing in the municipality on January 1 of each year; 2. aged 25–59; 3. male; 4. labor earnings above the minimum threshold, defined as one quarter of full-time work (13 weeks at 40 hours per week) at the lowest minimum wage across all prefectures; 5. primary source of income is employment income. For the household-level analysis, we restrict our sample to households that maintain the same household head throughout the panel period. This approach mitigates the impact of compositional changes when household heads relocate due to job transfers. Our final household-level sample consists of households with stable male heads aged 25-59, whose labor earnings exceed the minimum threshold defined above and whose primary income source is employment income. To estimate the ABB model (2), we further restrict the sample to balanced panels. Appendix A describes the construction of after-tax income, the household panel, and the sample restrictions in greater detail.

We do not impose the balanced panel restriction in our analysis of cross-sectional moments in Section 4; we refer to this sample as the cross-section sample. In contrast, we impose the balanced panel restriction for our analysis of idiosyncratic earnings risk in Sections 5 and 6, referring to this as the estimation sample.

Table 1 presents descriptive statistics of the cross-section sample, which includes both individual-level and household-level data. We focus on the cross-section sample statistics in the main text as it provides a more representative view of our selected sample, while the balanced panel restriction used for the estimation sample is primarily a technical requirement for our econometric analysis. To save space, we present the descriptive statistics of the estimation sample in the appendix.

The individual-level sample comprises 141,211 observations of prime-age males with an average age of 41.4 years. Average annual labor earnings are approximately 7.0 million yen, with substantial variation (standard deviation of 4.8 million yen) and a right-skewed distribution, as evidenced by the median (6.2 million yen) being lower than the mean. After-tax income averages 5.3 million yen, reflecting the progressive nature of the tax system.

The household-level sample contains 118,313 observations with male household heads averaging 42.0 years of age. Households have an average size of 2.94 members, with household labor earnings averaging 8.8 million yenhigher than individual earnings due to potential secondary earners. When equivalized by the square root of household size to account for economies of scale, average household income becomes 5.4 million yen, comparable to individual-level figures. The income distribution exhibits similar patterns of right skewness at both individual and household levels, with the 90th percentile substantially higher than the median across all income measures.

Table 1: Descriptive Statistics of the Cross-Section Sample

	Obs	Mean	Std. Dev.	P10	P50	P90
Individual-level sample						
Age	141211	41.36	9.04	29	41	54
Labor earnings	141211	6989.97	4763.52	3069.34	6223.65	11247.39
Total income	141211	7031.69	4884.75	3083.80	6238.59	11310.06
After-tax income	141211	5303.34	2996.22	2504.97	4889.44	8294.30
Household-level sample						
Household head's age	118313	42.03	8.83	30	42	54
Household size	118313	2.94	1.30	1.00	3.00	4.00
Labor earnings	118313	8775.24	5524.69	3930.00	7838.32	14266.56
Total income	118313	8857.85	5727.46	3957.35	7879.89	14402.94
After-tax income	118313	6764.15	3714.65	3198.87	6168.38	10793.79
Labor earnings (equivalized)	118313	5311.83	3153.21	2694.42	4747.72	8301.04
Total income (equivalized)	118313	5356.72	3236.77	2717.18	4774.18	8365.85
After-tax income (equivalized)	118313	4097.28	2071.23	2201.40	3732.61	6286.88

Notes. Columns designated as Obs, Mean, Std. Dev., P10, P50, and P90 report the number of observations, mean, standard deviation, and the 10th, 50th (median), and 90th percentiles of the respective variable. The monetary values are deflated by the Core CPI (excluding fresh food, 2020 base year) and expressed in thousands of yen. For the household-level sample, the equivalized income is calculated by dividing the household income by the square root of household size.

3.1 Distribution of Idiosyncratic Labor Earnings

This section documents the distribution of log labor earnings and its growth rate, as well as the variation in the idiosyncratic component of log earnings relative to overall log earnings.

Figure 1 presents histograms of log labor earnings and their growth rates at both individual and household levels. Panels (a) and (b) show individual-level data, while (c) and (d) show household-level data. Panel (a) shows that the distribution of individual log labor earnings is left-skewed with a longer tail on the left side. Panel (b) reveals that the distribution of changes in log labor earnings exhibits extremely high kurtosis, indicating a significant departure from normality. Panels (c) and (d) display similar distributions at the household level, with patterns closely resembling those observed in the individual-level data.

Figure 2 illustrates the age profile of log labor earnings and earnings growth. Panel (a) shows the coefficients on age dummies from regression (1), demonstrating that individual earnings follow an increasing and concave pattern with age. Panel (b) presents box plots of changes in log labor earnings by age, revealing that earnings growth rates decrease with age, reflecting the inverse U-shape of the age profile, Notably, the dispersion of growth rates is substantially larger for younger males (ages 25-26) compared to older individuals. The dispersion becomes smaller and more stable for individuals aged 27-54, partially supporting the empirical specification of a homogeneous lifecycle profile in equation (1).

At the household level, Panel (c) shows that the age effects of log equivalized household labor earnings (referenced to household head's age) also increase with age but follow a slightly different pattern. Unlike the individual-level data, the household age profile is slightly convex for ages above 27, possibly reflecting changes in sample composition with an increasing share of married couples and families with children. Panel (d) shows that household earnings growth rates remain relatively stable across ages, with slightly higher values for households with heads aged 25-26. The dispersion of household growth rates is higher than that of individual-level data, with households headed by individuals in their 40s exhibiting smaller dispersion.

Figure 3 compares the variances of log labor earnings and residuals, emphasizing the explanatory power of residual variations. Panels (a) and (b) show individual-level data, while (c) and (d) show household-level data. Panel (a) reveals that the variation in log earnings residuals is much smaller than that of log labor earnings. Panel (b) shows that the dispersion of log labor earnings can be mostly explained when individual fixed effects (μ_i) are added to the log earnings residuals. Panels (c) and (d) show similar patterns for household-level data, though the variation in log earnings residuals appears to explain a larger share of the variation in log labor earnings compared to individual-level data.

Table 2 quantifies the share of residual income in the variance of log measures, showing the ratio of the variance of log earnings residuals to the variance of log labor earnings, and corresponding ratios for total income and after-tax income. This share is 12–14% for individual-level data, while it is 15–18% for household-level data.

Figure 1: Histogram of Log Labor Earnings

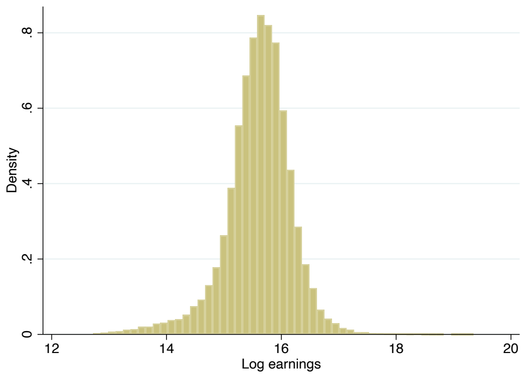

(a) Log Earnings, Male

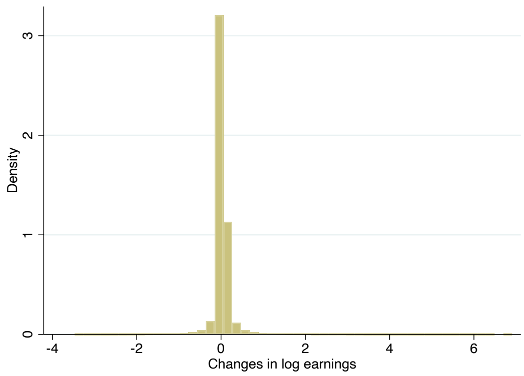

(b) Earnings Growth, Male

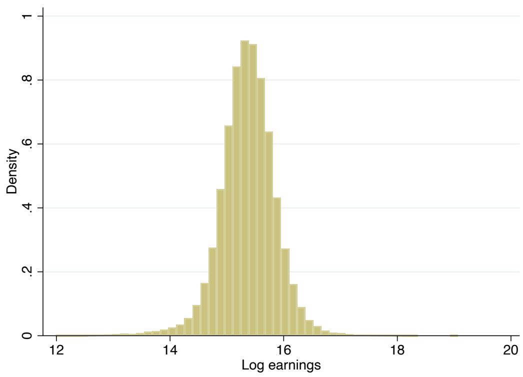

(c) Log Earnings, Household

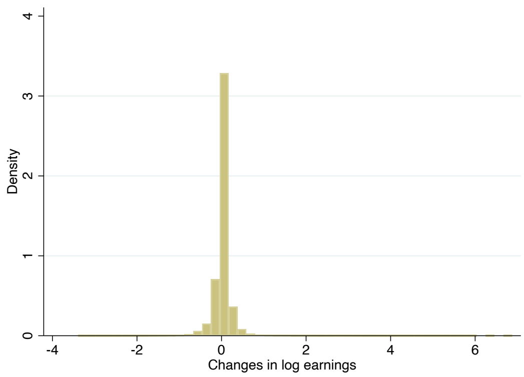

(d) Earnings Growth, Household

Notes. Panels (a) and (b) show individual-level data, while (c) and (d) show household-level data.

Figure 2: Age Profile of Log Labor Earnings and Earnings Growth

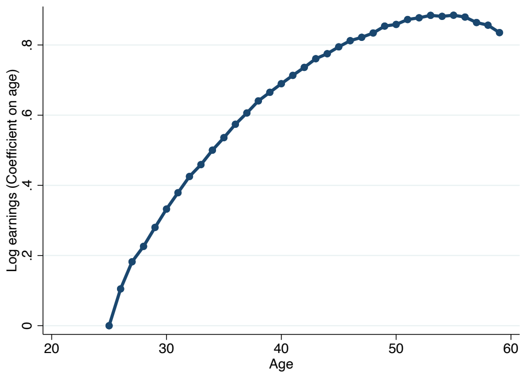

(a) Log Earnings, Male

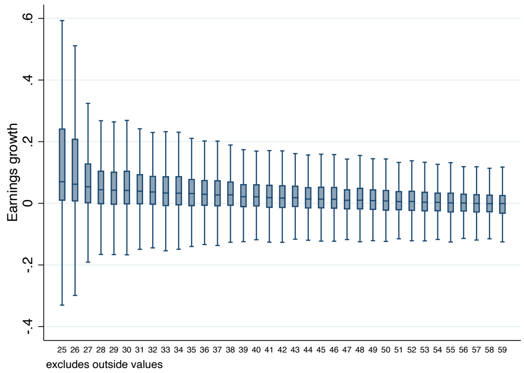

(b) Earnings Growth, Male

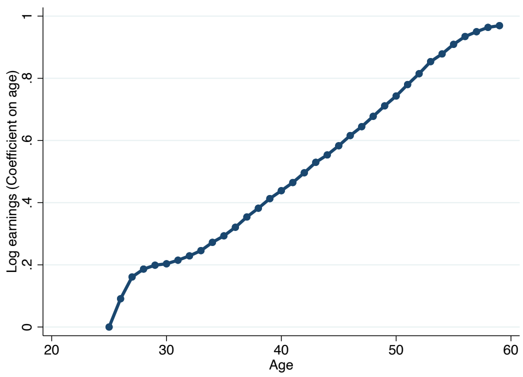

(c) Log Earnings, Household

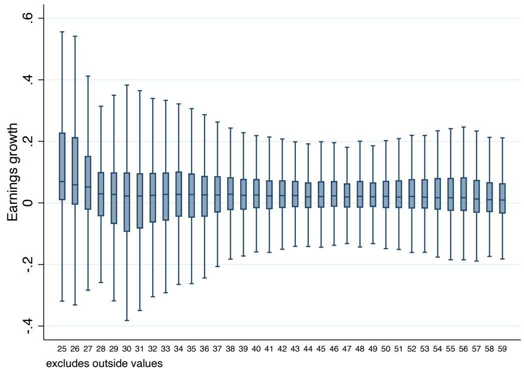

(d) Earnings Growth, Household

Notes. Panels (a) and (b) show individual-level data, while (c) and (d) show household-level data.

Figure 3: Log Labor Earnings and Residuals

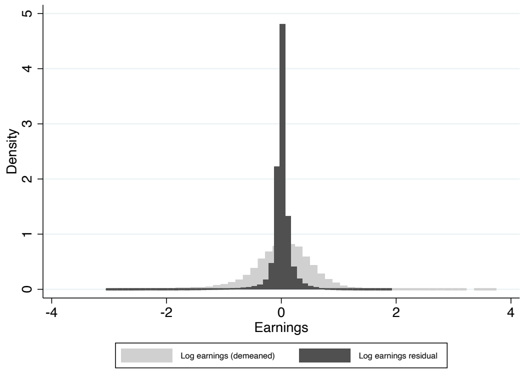

(a) y_{it}, Male

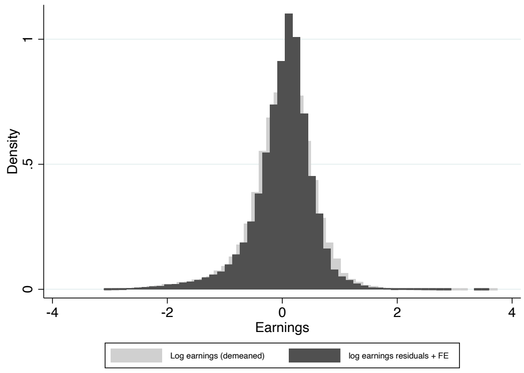

(b) $\mu_i + y_{it}$, Male

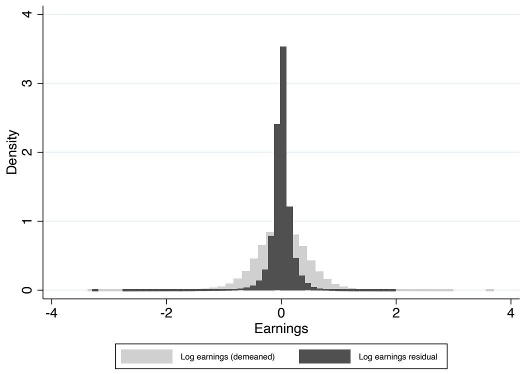

(c) y_{it}, Household

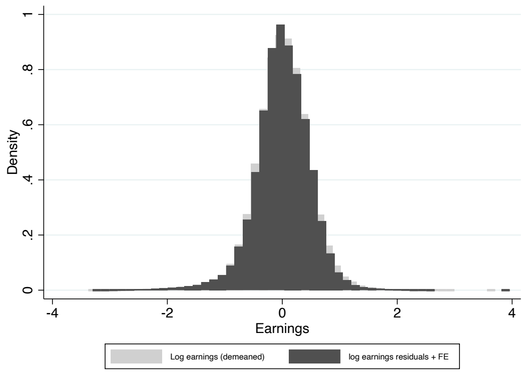

(d) $\mu_i + y_{it}$, Household

Notes. Panels (a) and (b) show individual-level data, while (c) and (d) show household-level data.

Table 2: Share of Residual Income in Variance of Log Measure

	Labor Earnings	Total income	After-tax income
Individual data			
$Var(\ln Y_{it})$	0.3423	0.3418	0.2959
$Var(y_{it)}$	0.0432	0.0416	0.0425
$Var(y_{it)}/Var(\ln Y_{it})$	0.1263	0.1217	0.1438
Household data			
$Var(\ln Y_{it})$	0.2460	0.2460	0.2146
$Var(y_{it)}$	0.0393	0.0385	0.0383
$Var(y_{it)}/Var(\ln Y_{it})$	0.1599	0.1567	0.1785

4 Income Inequality over Time and over Life-Cycle

Before examining the distribution of idiosyncratic components of log earnings and income, we first document the empirical patterns of inequality measures for the raw earnings and income data in the cross-section sample. All data are deflated by the Core CPI, and household aggregates are equivalized by the square root of household size. Figure 4 presents the level and changes in inequality measures for our individual-level sample from 2011 to 2021. Panels (a), (b), (c), and (d) show the variance of logarithm, the Gini coefficient, the ratio of the 50th to the 10th percentile, and the ratio of the 90th to the 50th percentile, respectively. In each panel, we compare the inequality measures of labor earnings, total income, and after-tax income.

The variances of log earnings and total income show a slight decline from 2011 to 2021, while the variance of log after-tax income shows a larger decline. Gini coefficients for labor earnings and total income remain relatively constant over time, while the Gini coefficient for after-tax income shows a decline. The ratio of the 50th to the 10th percentile declines for all three income measures, while the ratio of the 90th to the 50th percentile shows a slight increase for labor earnings and total income but remains stable or slightly declines for after-tax income. In all income inequality measures, after-tax income shows the lowest level of inequality compared to labor earnings and total income, highlighting the redistributive role of Japan's tax system.

Figure 5 presents the same inequality measures for our household-level sample. In this figure, we compare the inequality measures of household head labor earnings, household labor earnings, household total income and household after-tax income.

A notable pattern emerging from this figure is that all inequality measures for *household head* earnings remain stable or increase over time, while those for *household* earnings and income show a decline. As in the individual-level data, the inequality measures of after-tax income are the lowest among the four income measures.

To further examine the role of Japan's income tax system in reducing the level of inequality, Figure 6 presents the same inequality measures for household-level total income, after-tax income and after-tax income without income deductions. With Gini coefficient and the ratio of the 90th to the 50th percentile, there is no significant difference between after-tax income and after-tax income without income deductions, with the latter lying slightly above the former since 2017, indicating some progressivity of the income deductions for those years.

The variance of log and the ratio of the 50th to the 10th percentile, however, exhibit an opposite pattern, with the after-tax income without deductions showing a lower level of inequality than after-tax income. This seems surprising at first, but it may reflect the fact that income deductions matter only when taxable income is greater than the eligible deduction amount and the deduction amount is fixed within income brackets. This means that income deductions are not effective in addressing inequality at the bottom of the income distribution. Note that due to data limitations, we cannot examine the effect of public transfers such as social security benefits, which are likely to be more important in reducing inequality measures at the bottom of the income distribution.

Figure 7 presents lifecycle inequality patterns using cross-sectional age profiles organized in 5-year groups from ages 25-29 to 55-59. The analysis addresses the fundamental identification problem of separating age, time, and cohort effects by employing two distinct approaches. Panel (a) and (c) control for time effects: the variance of logarithm for age group a at time t, denoted $M(a,t)$, is regressed on a full set of age group and year dummies, with the age profile computed from the predicted age-portion of the age-time regression. Panels (b) and (d) control for cohort effects, omitting time effects. At the individual level, earnings inequality exhibits a clear monotonic increase with age, rising from approximately 0.05 at ages 25-29 to 0.15 at ages 55-60, consistent with the accumulation of permanent wage shocks over the working life. Total income inequality follows a similar but moderately attenuated pattern, while after-tax income inequality remains relatively compressed around 0.05-0.08 across all age groups, reflecting the redistributive effects of the tax system.

The household-level analysis reveals markedly different dynamics, with total income and after-tax income displaying distinctive J-curve or U-curve patterns that initially decline or remain flat during early career years before increasing in later working years. This contrasts sharply with head earnings and total earnings, which maintain upward trajectories similar to individual measures. Note that household head labor earnings are presented in raw form without equivalization, while household earnings, total income, and after-tax income are equivalized. Therefore, the variance of log household head earnings reflects earnings variation in labor markets, while the variance of log equivalized household earnings and income measures dispersion in household living standards. The J-curve pattern in equivalized household income reflects the interaction of permanent wage shocks, endogenous labor supply responses, household insurance mechanisms including spousal adjustments and income pooling, and progressive income taxation. These findings suggest that household formation and policy interventions substantially alter inequality dynamics compared to individual-level measures, providing important smoothing during early and middle career years when individual earnings inequality is rising most rapidly.

Figure 8 presents the effective tax rate, which is defined by the ratio of the amount of income tax paid to total income, and the combined income tax and social security contribution rate against total income. It shows that the effective tax rate and the combined rate, not surprisingly, increase with household total income, although they have significant variation beyond the income level.

To further examine the heterogeneity in the effective tax rate and the combined tax and social security contribution rate, we regress those rates on household type dum-

Figure 4: Inequality Measures

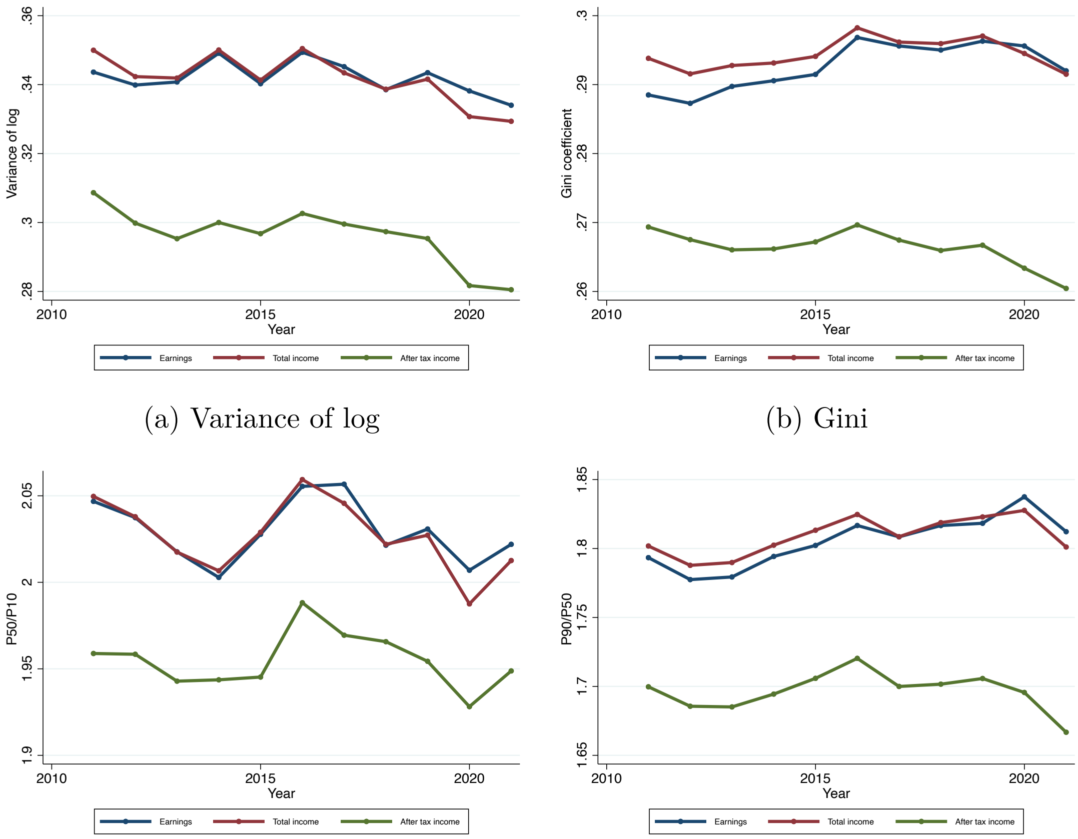

(a) Variance of log

(b) Gini

(c) P50/P10

(d) P90/P50

Notes. Individual-level data are used.

mies and a fourth-order polynomial of orthogonalized household total income. Table 3 reports the estimation results. After controlling for household total income level, households with household head and spouse both working, with or without children, have the lowest effective tax rate, which is around 4% lower than single households. Because income taxes are levied on individual income, lower progressive tax rates apply to head and spouse's income separately, compared to the single earner's income, conditional on the same household total income. It appears that dual-earner households pay higher social security contributions compared to one-earner households, resulting in a smaller gap between the household types in the combined tax and social security contribution rate. There seems to be a significant difference of around 0.6% between married, one working, no kids and married, one working, with kids households.

Figure 5: Inequality Measures

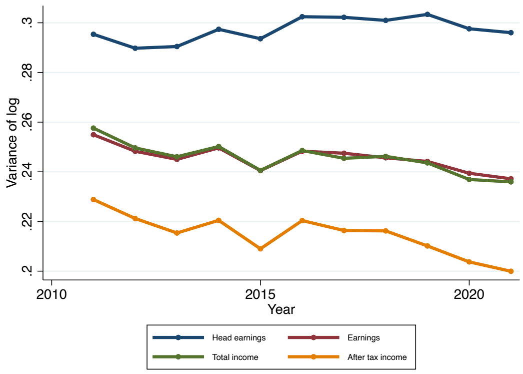

(a) Variance of log

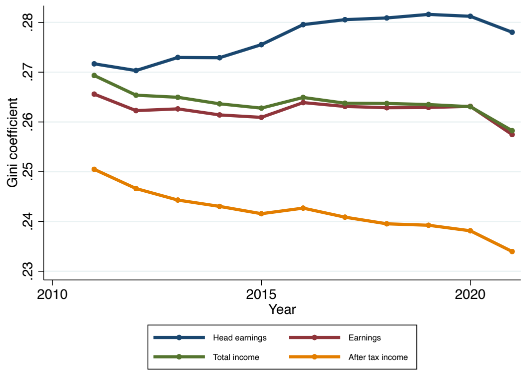

(b) Gini

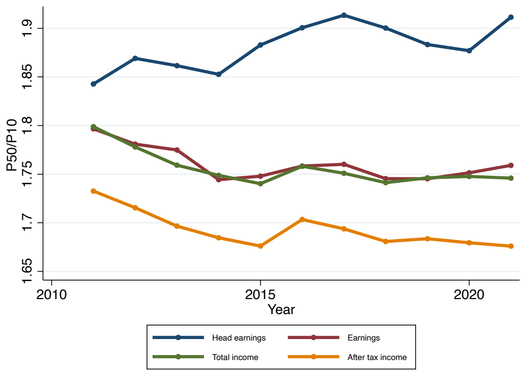

(c) P50/P10

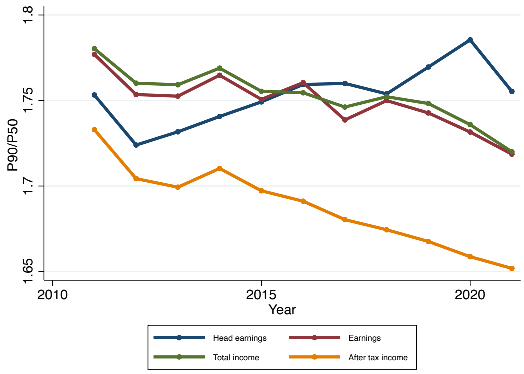

(d) P90/P50

Notes. Household-level data are used.

Figure 6: Role of Income Deductions

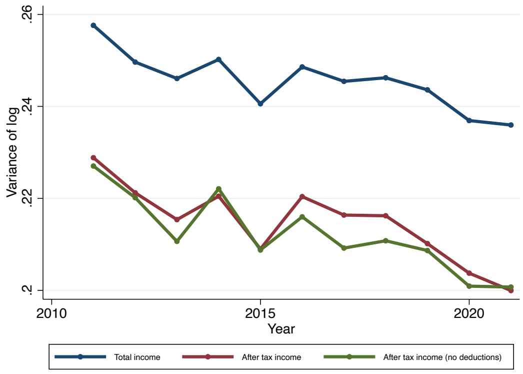

(a) Variance of log

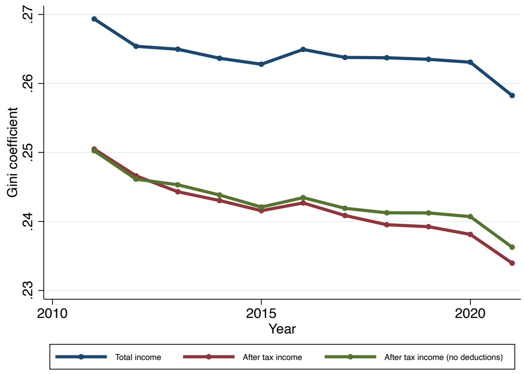

(b) Gini

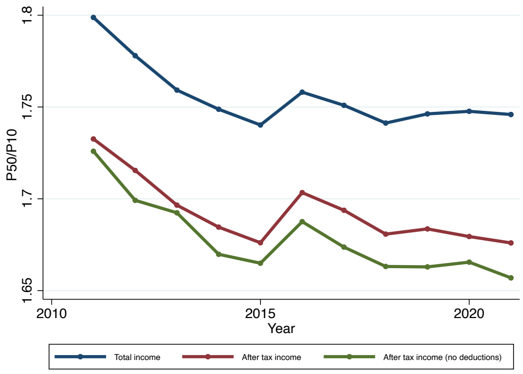

(c) P50/P10

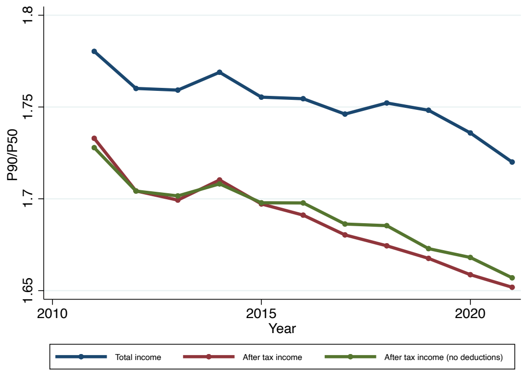

(d) P90/P50

Notes. Household-level data are used.

Figure 7: Lifecycle Inequality

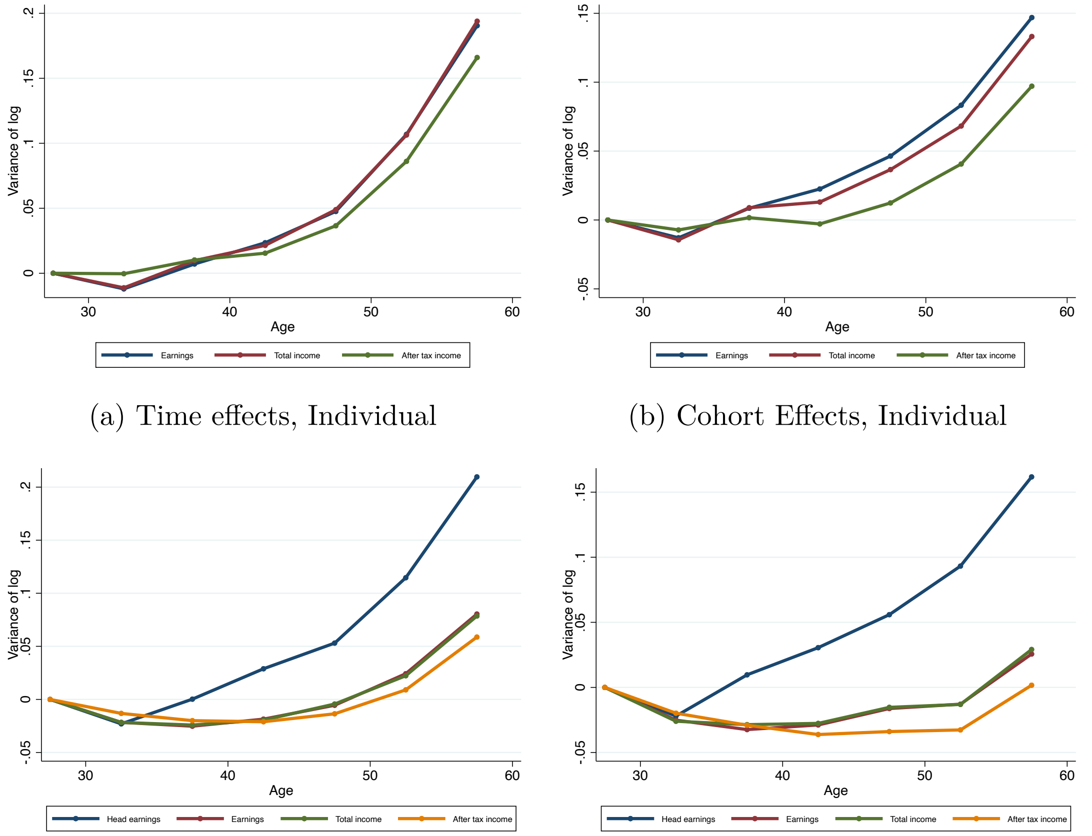

(a) Time effects, Individual (b) Cohort Effects, Individual

(c) Time Effects, Household (d) Cohort Effects, Household

Notes. Panels (a) and (b) show individual-level data, while (c) and (d) show household-level equivalized data.

Figure 8: Effective Tax and Social Security Contribution Rate

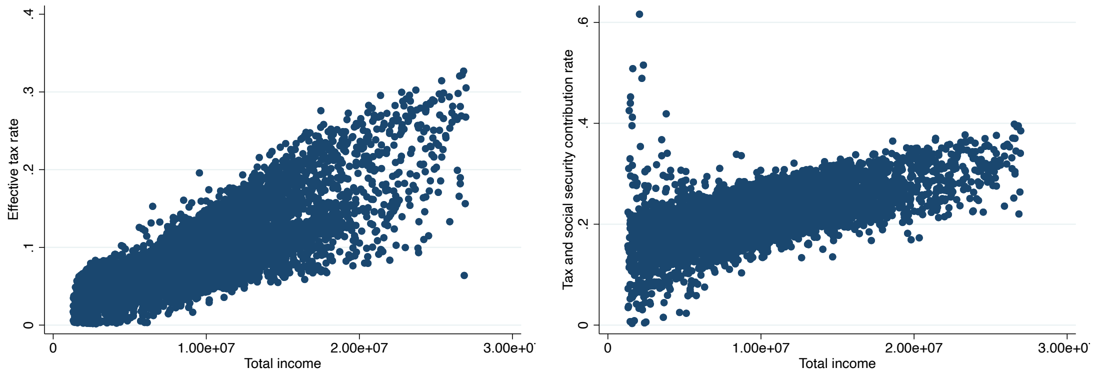

(a) Effective Tax Rate (b) Tax and Social Security Rate

Notes. The left panel presents a scatter plot of effective income tax rates against household total income, while the right panel displays a scatter plot of combined income tax and social security contribution rates against household total income. The figure uses household-level data from a municipality covering the years 2011 to 2021. A 10% random subsample is employed, and observations are restricted to the interquartile range between the 1st and 99th percentiles of household total income to ensure data confidentiality.

Table 3: Effect of Household Type on Effective Tax And Social Security Contribution Rate

	(1) Effective Tax Rate	(2) Tax plus Social Security
Single, with kids	-0.0210***	-0.0172***
	(0.000810)	(0.00123)
Married, both working, no kids	-0.0446***	-0.0313***
	(0.000285)	(0.000434)
Married, both working, with kids	-0.0436***	-0.0328***
	(0.000256)	(0.000390)
Married, one working, no kids	-0.0256***	-0.0222***
	(0.000265)	(0.000403)
Married, one working, with kids	-0.0173***	-0.0148***
	(0.000194)	(0.000295)
Income Controls	✓	✓
Observations	118,320	118,313
R-squared	0.758	0.445
Adj. R-squared	0.758	0.445

Notes. Household-level data are used. The dependent variable is the effective tax rate in Column (1) and the effective tax plus social security contribution rate in Column (2). A fourth-order polynomial of orthogonalized income is included as a control. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5 Quantile-based Conditional Moments of Idiosyncratic Earnings and Income

This section examines the distribution of the idiosyncratic components of log earnings and income, which are residual y_{it} from regression (1), by estimating quantile functions. We characterize the conditional distribution of the idiosyncratic component y_{it} through a quantile function Q_y as follows:

$$y_{it} = Q_y(y_{it-1}, age_{it}, u_{it}), \quad (u_{it}|y_{it-1}) \sim \text{iid Uniform}(0,1).$$

By estimating Q_y using quantile regression methods, we can calculate conditional moments of y_{it} based on Q_y. First, conditional persistence $\rho(y_{it-1}, age_{it}, \tau)$ is defined as follows:

$$\rho(y_{it-1}, age_{it}, \tau) := \frac{\partial Q_y(y_{it-1}, age_{it}, \tau)}{\partial y}.$$

The conditional dispersion measure $\sigma(y_{it-1}, age_{it}, \tau)$ for $\tau \in (1/2, 1)$ is

$$\sigma(y_{it-1}, age_{it}, \tau) := Q_y(y_{it-1}, age_{it}, \tau) - Q_y(y_{it-1}, age_{it}, 1-\tau).$$

Conditional skewness $sk(y_{it-1}, age_{it}, \tau)$ for $\tau \in (1/2, 1)$ is

$$sk(y_{it-1}, age_{it}, \tau) := \frac{Q_y(y_{it-1}, age_{it}, \tau) + Q_y(y_{it-1}, age_{it}, 1-\tau) - 2Q_y(y_{it-1}, age_{it}, \frac{1}{2})}{Q_y(y_{it-1}, age_{it}, \tau) - Q_y(y_{it-1}, age_{it}, 1-\tau)}.$$

Conditional kurtosis $kur(y_{it-1}, age_{it}, \tau, \alpha)$ for $\tau \in (1/2, 1)$ and $\alpha < 1-\tau$ is

$$kur(y_{it-1}, age_{it}, \tau, \alpha) := \frac{Q_y(y_{it-1}, age_{it}, 1-\alpha) - Q_y(y_{it-1}, age_{it}, \alpha)}{Q_y(y_{it-1}, age_{it}, \tau) - Q_y(y_{it-1}, age_{it}, 1-\tau)}.$$

For estimation, we specify Q_y as follows:

$$Q_y(y_{it-1}, age_{it}, \tau) = \sum_{k=0}^{K}\sum_{j=0}^{J} a_{k,j}(\tau)\varphi_k(y_{it-1})\psi_j(age_{it}),$$

where $a_{k,j}(\tau)$ represents parameters, $\varphi_k(y_{it-1})$ represents Hermite polynomials of y_{it-1}, and $\psi_j(age_{it})$ represents Hermite polynomials of age_{it}.

5.1 Individual-level Labor Earnings

This section presents the quantile-based conditional moments of idiosyncratic earnings y_{it} using our individual-level sample. The conditional persistence $\rho(y_{it-1}, \tau)$ exhibits substantial heterogeneity across the distribution of lagged earnings residuals and the magnitude of shocks as shown in Figure 9. Persistence reaches its highest levels, approximately

0.8–0.9, for two distinct groups: individuals with very low lagged earnings experiencing large negative shocks (low current percentile) and those with very high lagged earnings experiencing large positive shocks (high current percentile). Conversely, persistence attains its minimum values of approximately 0.2 for individuals at opposite extremes of the shock distribution—those with very low lagged earnings receiving large positive shocks and those with very high lagged earnings experiencing large negative shocks. This pattern of nonlinear persistence, where the influence of past earnings is weakest for large 'reversal' shocks, is strikingly similar to the findings of Arellano et al. (2017) using PSID data for the United States.

Economically, this nonlinearity suggests that while earnings paths are quite stable for individuals on a consistent trajectory, large shocks can act as 'career resets.' For instance, a large negative shock to a high-earner may represent a job displacement or a significant health event that makes their past earnings history a poor predictor of future income. Conversely, a large positive shock to a low-earner could signify a major promotion or career change that sets them on a new, higher trajectory.

Figure 10 shows that conditional persistence generally increases with age. Workers aged 35, 45, and 55 years exhibit patterns of conditional persistence variation across lagged earnings levels and shock magnitudes that closely mirror the pooled sample. However, the 25-year-old cohort displays a distinctly different pattern, characterized by persistence that decreases with the magnitude of shocks, diverging from the established pattern observed in older age groups. This age-related heterogeneity may reflect differences in career stability and labor market attachment across the life cycle.

Figure 11 presents the conditional dispersion $\sigma(y_{it-1}, \tau)$ against lagged earnings residuals. Panel (a) shows that conditional dispersion averages approximately 0.16 and exhibits a pronounced U-shaped pattern across the lagged earnings distribution, reaching peak values of approximately 0.2 at both tails of the distribution. This pattern indicates that earnings volatility is highest for individuals at the extremes of the earnings distribution. Panel (b) presents age-specific patterns and reveals that conditional dispersion decreases with age. Workers aged 35, 45, and 55 exhibit conditional dispersion patterns consistent with the aggregate sample, with dispersion concentrated at the tails of the lagged earnings distribution. The 25-year-old cohort constitutes an exception, displaying the highest overall dispersion levels that decrease monotonically with lagged earnings levels, suggesting greater earnings instability early in careers.

Figure 12 reports the conditional skewness $sk(y_{it-1}, \tau)$. Panel (a) demonstrates that conditional skewness decreases monotonically with lagged earnings levels in the whole sample. Skewness assumes positive values for low earners and transitions to negative values for high earners, indicating that individuals with low earnings face positively skewed earnings shocks (greater probability of large positive changes) while those with high earnings encounter negatively skewed shocks (greater probability of large negative changes).

Panel (b) shows that this decreasing relationship between skewness and lagged earnings remains consistent across most age groups. A notable exception occurs among 25-year-old workers, who exhibit positive skewness across all levels of lagged earnings with a mildly concave profile characterized by reduced skewness at both extremes of the lagged earnings distribution. Additionally, skewness tends to decrease with age, with the 25-year-old cohort again representing a departure from this general pattern.

Figure 13 reveals that conditional kurtosis $kur(y_{it-1}, \tau, \alpha)$ consistently exceeds levels expected under a normal distribution, indicating the presence of fat tails in the earnings shock distribution. Panel (a) demonstrates that kurtosis exhibits an increasing relationship with lagged earnings levels, suggesting that high earners face more extreme earnings shocks than low earners. Panel (b) shows that kurtosis patterns remain broadly similar across age groups, with all cohorts displaying increasing kurtosis relative to lagged earnings. Kurtosis demonstrates an increasing trend with age, though the profiles of conditional kurtosis against lagged earnings exhibit crossing patterns at the extremes of the distribution, indicating complex age-related heterogeneity in the tail behavior of earnings shocks.

Figure 9: Persistence of Log Earnings Residuals

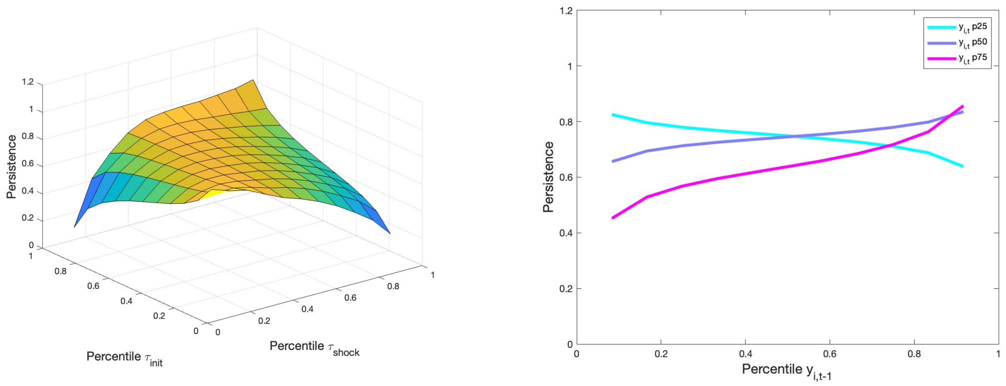

Notes. Individual-level data for 2011-2021 for a municipality.

5.2 Household-level Earnings and Income

Our analysis of household-level data reveals several important patterns in the conditional moments of log earnings residuals. The results are presented in Figures 14 through 18 and demonstrate both similarities and notable differences compared to individual-level findings.

The persistence of household labor earnings residuals exhibits a pattern broadly consistent with individual-level data, as shown in Figures 14–15. In contrast to the individual-level data, however, Figure 15 shows that households with 25-year-old heads display a persistence pattern similar to both the overall sample and other age groups.

Figure 10: Persistence of Log Earnings Residuals By Age

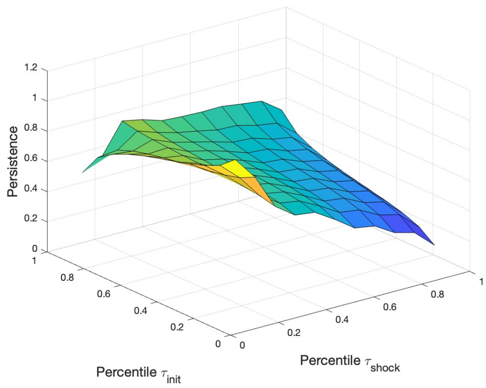

(a) 25 years old

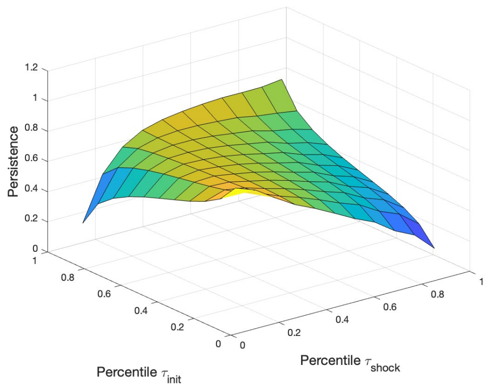

(b) 35 years old

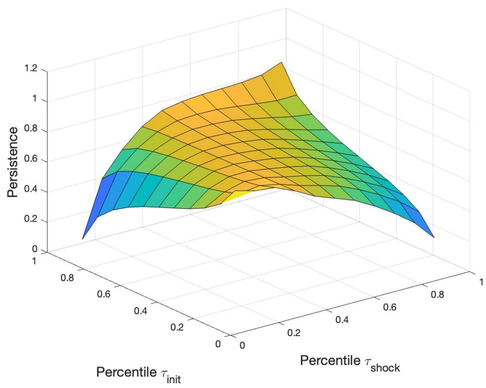

(c) 45 years old

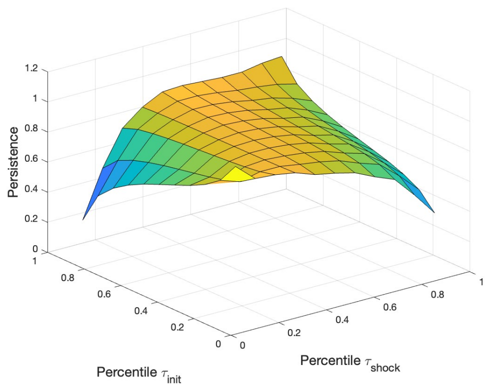

(d) 55 years old

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 11: Conditional Dispersion of Log Earnings Residuals

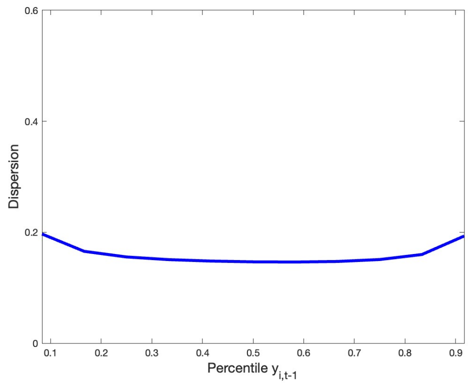

(a) Whole Sample

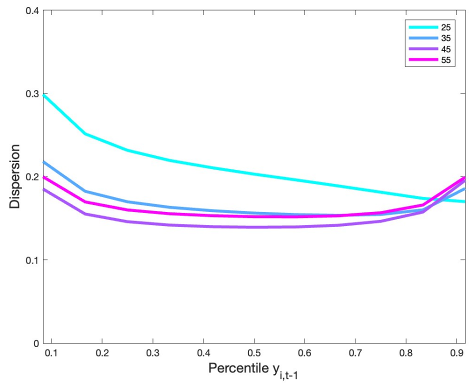

(b) Age Group

Notes. Individual-level data for 2011-2021 for a municipality.

The conditional dispersion of household labor earnings residuals presents several noteworthy characteristics. First, the dispersion levels are systematically higher than those observed in individual-level data. Both the aggregate sample and age-specific subsam-

Figure 12: Conditional Skewness of Log Earnings Residuals

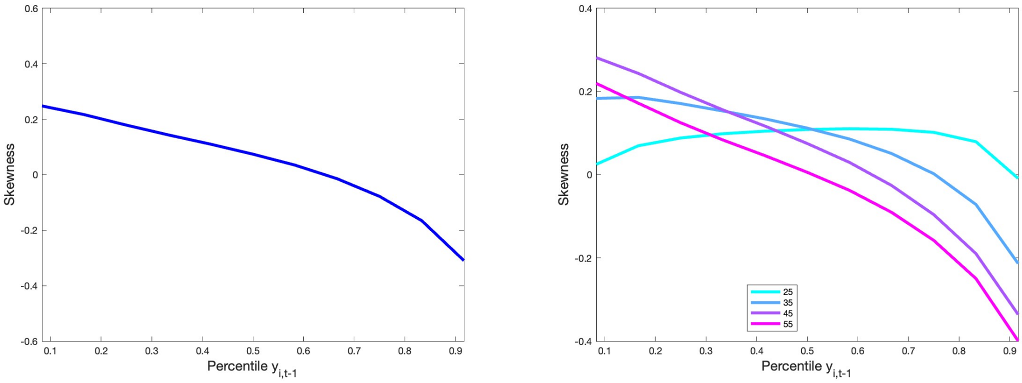

(a) Whole Sample (b) Age Group

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 13: Conditional Kurtosis of Log Earnings Residuals

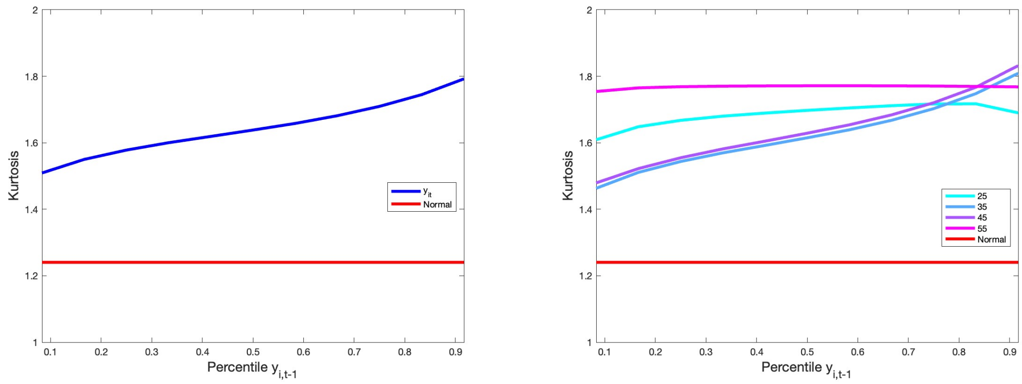

(a) Whole Sample (b) Age Group

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 14: Persistence of Log Earnings Residuals

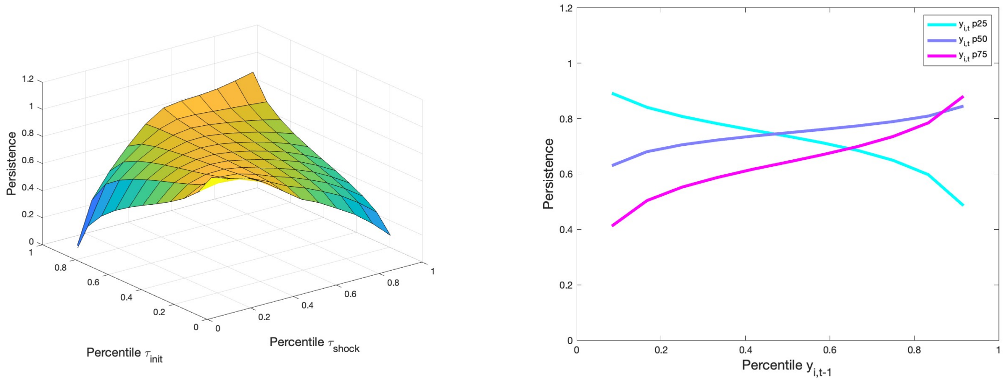

Notes. Household-level data for 2011-2021 for a municipality.

ples exhibit a pronounced U-shaped pattern across the lagged earnings distribution, as illustrated in Figure 16.

Age-related patterns in conditional dispersion show a general declining trend with age, with one notable exception: the 55-year-old cohort maintains dispersion levels comparable to those of the 35-year-old group, deviating from the otherwise monotonic age-related decline.

The conditional skewness patterns largely mirror those found in individual-level analysis. Figure 17 demonstrates a consistent negative relationship between conditional skewness and lagged earnings levels. Low-earning households face positively skewed earnings residuals, while high-earning households encounter negatively skewed distributions.

A key distinction from individual-level data emerges in the age-related patterns: skewness increases with age while maintaining the negative relationship with lagged earnings levels observed in the whole sample. This age-related increase in skewness represents a departure from individual-level findings.

Consistent with individual-level results, conditional kurtosis remains systematically higher than levels expected under a normal distribution assumption, with a positive relationship to lagged earnings levels, as shown in Figure 18.

The age-specific patterns of conditional kurtosis exhibit considerable heterogeneity without a clear ordering across age groups. The 45- and 55-year-old cohorts demonstrate increasing conditional kurtosis with lagged earnings levels. In contrast, the 25-year-old group exhibits a declining pattern, while the 35-year-old group displays an inverse U-shaped relationship between conditional kurtosis and lagged earnings levels.

Figure 15: Persistence of Log Earnings Residuals By Age

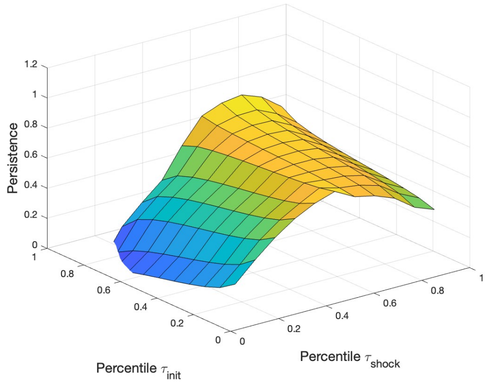

(a) 25 years old

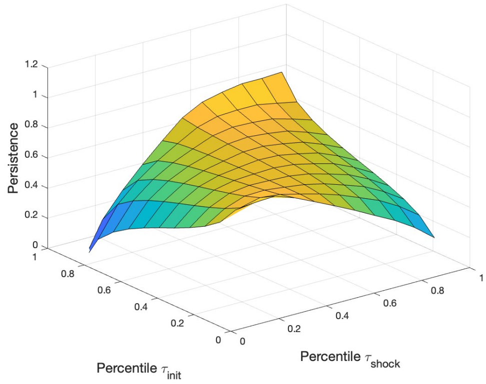

(b) 35 years old

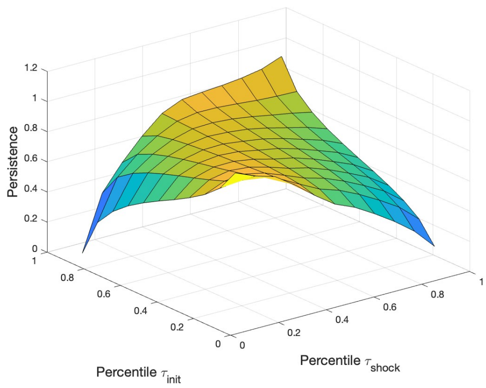

(c) 45 years old

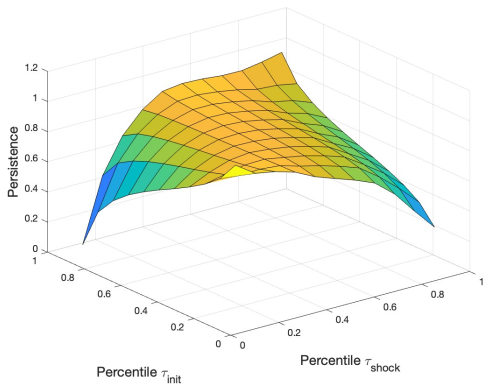

(d) 55 years old

Notes. Household-level data for 2011-2021 for a municipality.

Figure 16: Conditional Dispersion of Log Earnings Residuals

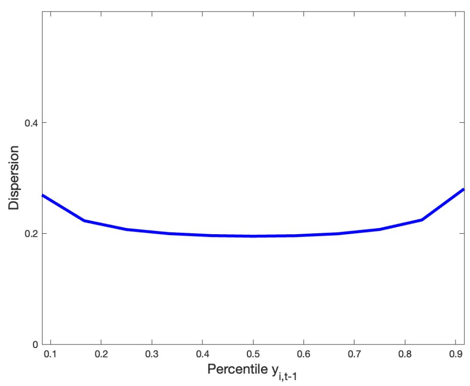

(a) Whole Sample

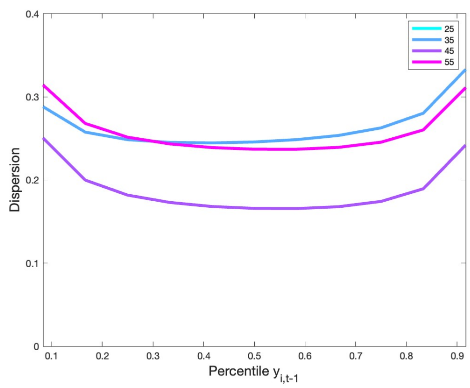

(b) Age Group

Notes. Household-level data for 2011-2021 for a municipality.

Figure 17: Conditional Skewness of Log Earnings Residuals

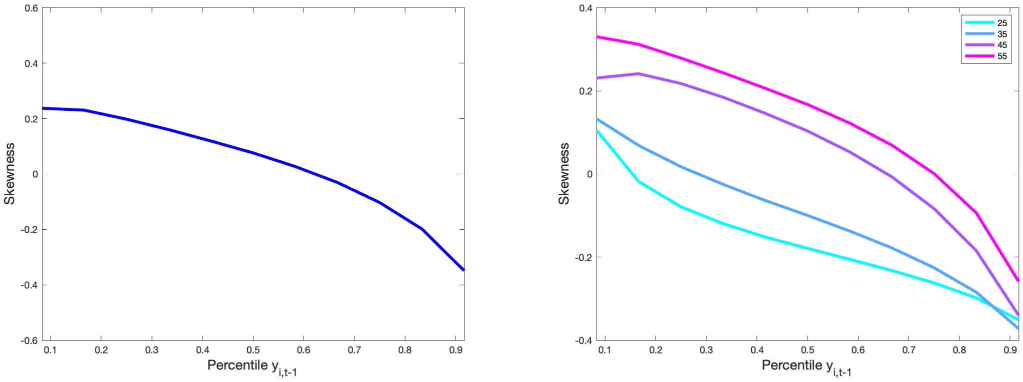

(a) Whole Sample (b) Age Group

Notes. Household-level data for 2011-2021 for a municipality.

Figure 18: Conditional Kurtosis of Log Earnings Residuals

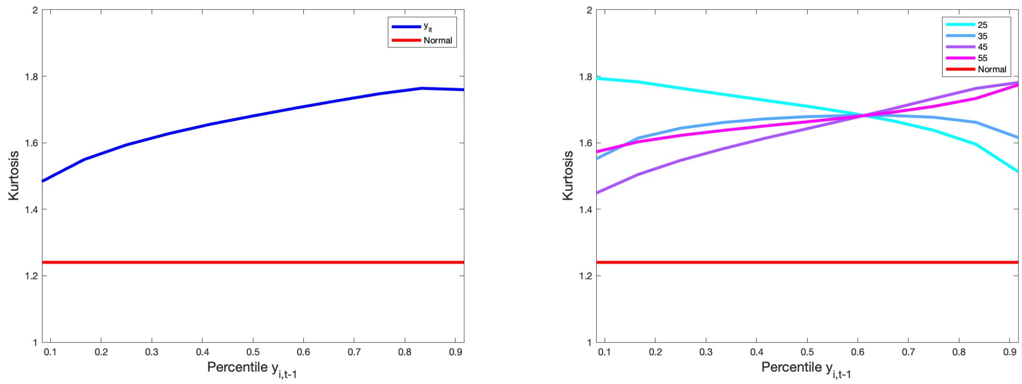

(a) Whole Sample (b) Age Group

Notes. Household-level data for 2011-2021 for a municipality.

6 Persistent and Transitory Shocks

Theories of consumption and savings emphasize the importance of distinguishing between persistent and transitory income shocks, as these have different implications for consumption smoothing and, consequently, welfare. Following ABB, we decompose y_{it} into a persistent component η_{it} and a transitory component ε_{it}, modeling these components as in (2). We re-state (2) here for clarity:

$$\begin{cases} y_{it} = \eta_{it} + \varepsilon_{it} \\ \eta_{it} = Q_{\eta}(\eta_{it-1}, age_{it}, u_{it}), \quad (u_{it}|\eta_{it-1}, \eta_{it-2}, \ldots) \sim \text{Uniform}(0,1) \\ \varepsilon_{it} \sim Q_{\varepsilon}(age_{it}, u_{it}^{\varepsilon}), \quad u_{it}^{\varepsilon} \sim \text{Uniform}(0,1) \\ \eta_{i1} \sim Q_{\eta_1}(age_{it},, u_{it}^{\eta_1}), \quad u_{it}^{\eta_1} \sim \text{Uniform}(0,1) \end{cases}$$

6.1 Estimation Results using Individual-level Earnings Data

In this section, we report estimation results using individual-level idiosyncratic earnings data. Given that persistent shocks have substantially larger effects on consumption and welfare compared to transitory shocks, we focus our analysis on the distribution of the persistent component η_{it}, which is presented in Figures 19 and 20. Conditional moments derived from Q_{ε} and Q_{η_1} are provided in the Appendix for completeness. To examine the implications of nonlinear persistence in η_{it} for idiosyncratic earnings dynamics, we calculate impulse responses to persistent shocks following the methodology of ABB, with results reported in Figures 23–24.

Figure 19 demonstrates that the conditional persistence derived from Q_{η} exhibits nonlinear relationships with both η_{it-1} and the magnitude of shocks. These patterns closely mirror those observed for y_{it}. Similarly, Figure 20 reveals age-specific patterns that are consistent with those presented in Figure 10.

To check how the estimated model captures the empirical patterns observed in the data more directly, we simulate the income data from the estimated model and compare the simulated data with the empirical data in terms of the cross-sectional and quantile-based moments.

Figure 21 compares the persistence estimates between empirical and simulated data. Panel (a) shows the empirical data patterns, while panel (b) displays the corresponding simulated data results. Panels (c) and (d) provide additional comparative views of the persistence estimates, demonstrating how well the estimated model replicates the observed persistence patterns in the actual data. The comparison reveals the model's ability to capture the nonlinear properties of income persistence across different specifications.

Figure 22 presents a comparison of income variance patterns across different age groups between the empirical and simulated datasets. The empirical data exhibits a distinctive U-

Figure 19: Persistence of η_{it}

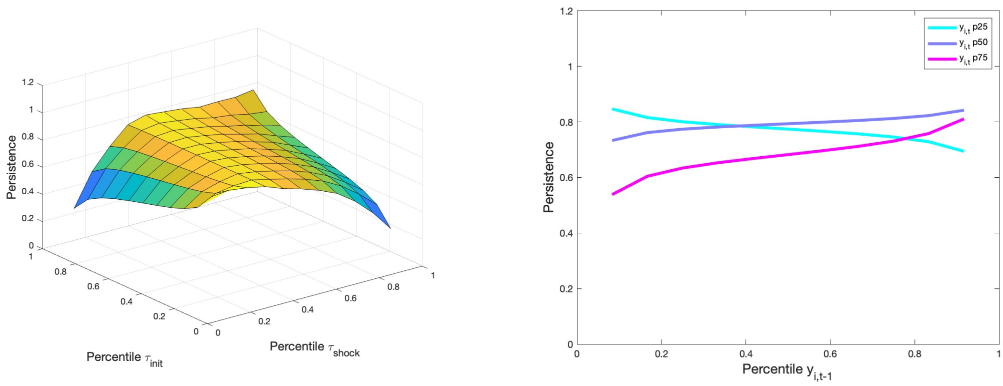

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 20: Persistence of η_{it} By Age

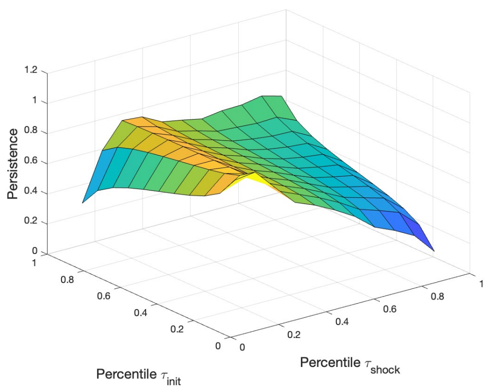

(a) 25 years old

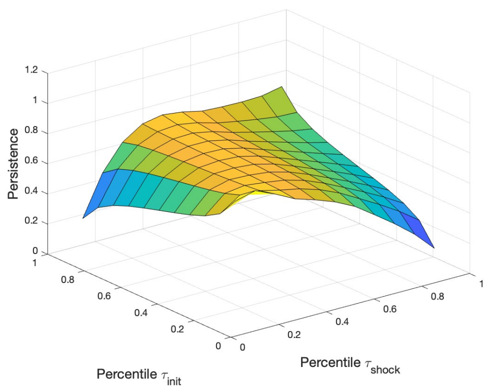

(b) 35 years old

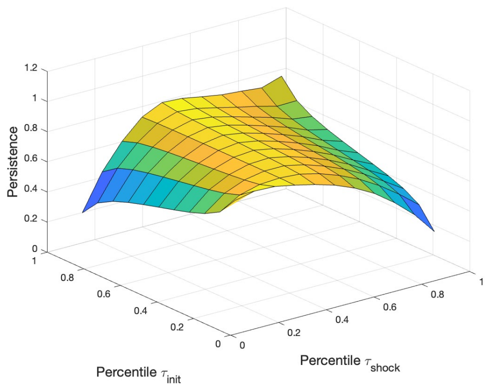

(c) 45 years old

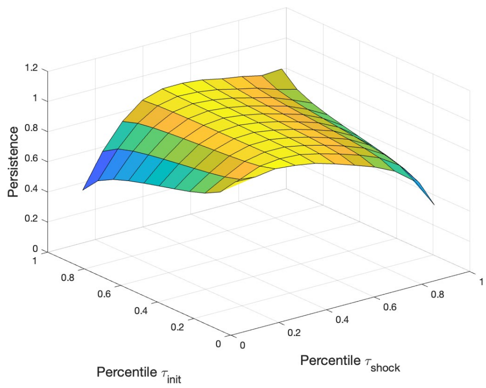

(d) 55 years old

Notes. Individual-level data for 2011-2021 for a municipality.

shaped pattern in income variance across age groups, with higher variance among younger workers, declining variance in middle age, and increasing variance again among older workers. This U-shaped pattern represents a key stylized fact that cannot be adequately explained by canonical linear models with constant shock variances. The estimated model demonstrates a strong ability to replicate this complex age-specific variance structure, successfully capturing the nonlinear life-cycle income dynamics and heterogeneity patterns across different demographic groups. While there are some gaps between the empirical and simulated series, the model effectively reproduces the key features of this U-shaped income variance pattern by age group.

Figure 21: Empirical vs. Simulated Data: Persistence Estimates

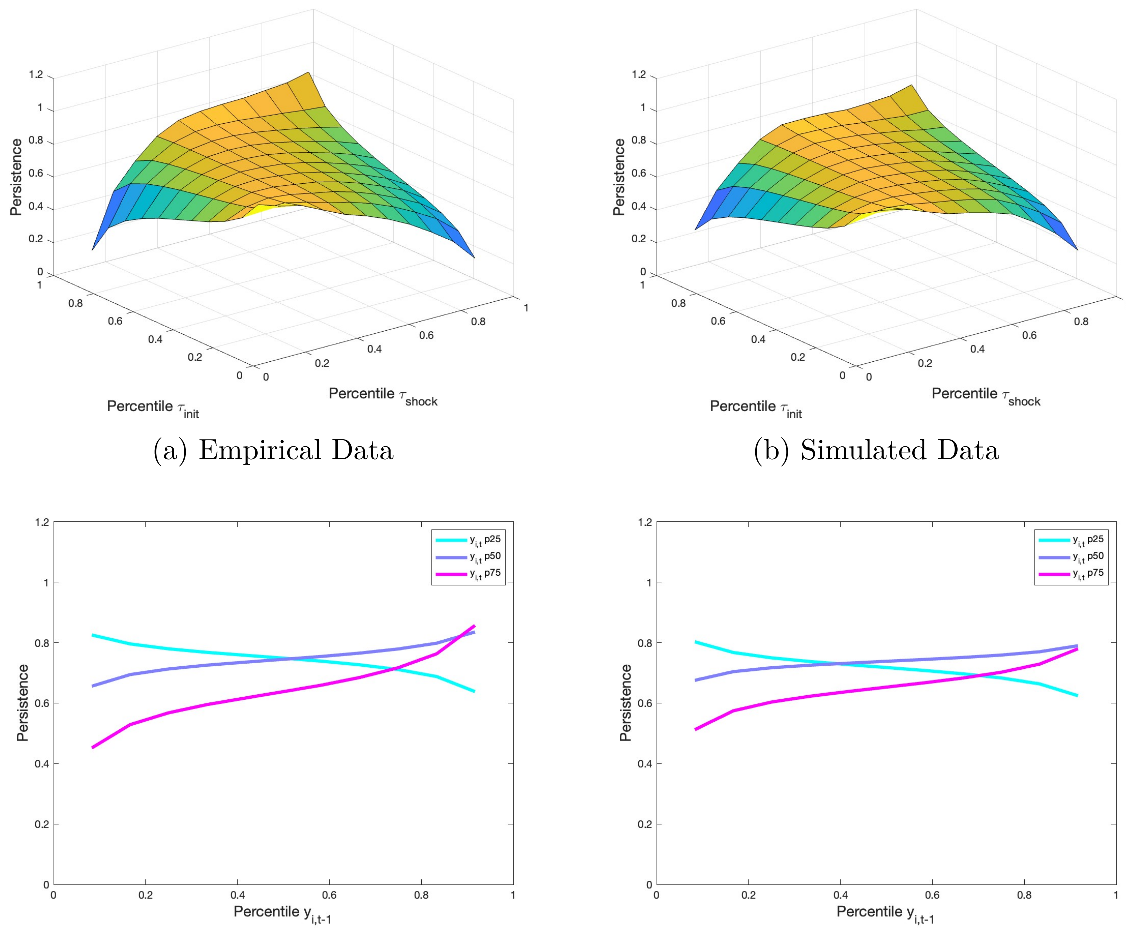

(a) Empirical Data (b) Simulated Data

(c) Empirical Data (d) Simulated Data

To analyze the dynamic effects of persistent earnings shocks, we follow the approach of ABB and construct impulse response functions through simulation. Our methodology proceeds in three steps: First, we simulate earnings trajectories $(y_{it})_{it}$ for individuals starting from a pre-specified initial quantile position τ_{init} at age 35. We generate two sets of trajectories: (i) a shock scenario $(y_{it}^{\text{shock}})_{it}$ where individuals receive a quantile shock τ_{shock} of either 0.1 or 0.9 at age 36, and (ii) a baseline scenario $(y_{it}^{\text{baseline}})_{it}$ where individuals

Figure 22: Empirical vs. Simulated Data: Variance by Age Group

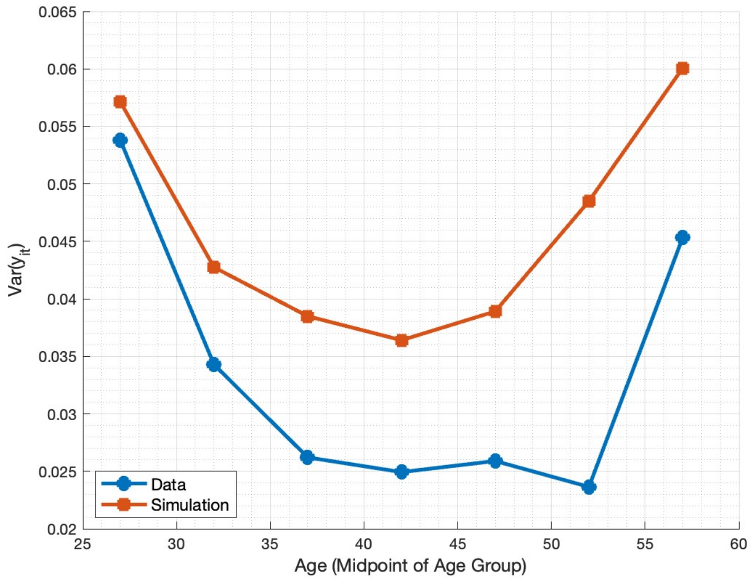

receive a neutral shock of $\tau_{\text{shock}} = 0.5$ at age 36.

Second, for each age t, we calculate the impulse response as the difference in median earnings between the shock and baseline scenarios:

$$IR(t) = \text{median}\left(y_{it}^{\text{shock}}\right) - \text{median}\left(y_{it}^{\text{baseline}}\right)$$

Third, we plot the impulse response function $IR(t)$ over the life cycle to visualize the persistence and evolution of earnings shocks.

Figure 23 presents the impulse response functions for different combinations of initial quantile positions (τ_{init}) and shock magnitudes (τ_{shock}). The results reveal substantial heterogeneity in shock persistence across the earnings distribution, reflecting the nonlinear nature of earnings dynamics captured in our model. For adverse shocks ($\tau_{\text{shock}} = 0.1$), the impact is notably larger for individuals initially positioned at higher quantiles ($\tau_{\text{init}} = 0.9$) compared to those at lower quantiles ($\tau_{\text{init}} = 0.1$). This asymmetric response occurs because persistence is lower for individuals with higher lagged earnings residuals. Conversely, for positive shocks ($\tau_{\text{shock}} = 0.9$), the impact is more pronounced for individuals initially at lower quantiles ($\tau_{\text{init}} = 0.1$) than for those at higher quantiles ($\tau_{\text{init}} = 0.9$), again reflecting the differential persistence patterns across the earnings distribution.

While our current framework does not permit estimation of welfare effects without modeling structural consumption responses, we can quantify the long-term impact of persistent earnings shocks by calculating their cumulative effects on lifetime earnings. Our calculation follows a three-step procedure. First, we use the simulated earnings trajectories from both shock and baseline scenarios. Second, we calculate discounted lifetime earnings for ages 35-59 for each individual:

$$\bar{Y}_i^{\text{shock}} = \sum_{t=35}^{59} \left(\frac{1}{1+r} \right)^{t-35} \exp(\hat{\beta}_t^{\text{age}}) \exp(y_{it}^{\text{shock}})$$

where r is the discount rate and $\hat{\beta}_t^{\text{age}}$ captures the estimated age profile of earnings. We calculate $\bar{Y}_i^{\text{baseline}}$ analogously for the baseline scenario. Third, we define the cumulative effect as the proportional difference in median lifetime earnings:

$$\text{Cumulative Effect} = \frac{\text{median}(\bar{Y}_i^{\text{shock}}) - \text{median}(\bar{Y}_i^{\text{baseline}})}{\text{median}(\bar{Y}_i^{\text{baseline}})}$$

Figure 24 displays these cumulative effects in a heatmap format, showing how the long-term impact of earnings shocks varies across initial quantile positions and shock magnitudes. The heterogeneous patterns observed in the impulse response functions translate into substantial variation in lifetime earnings effects. For adverse shocks ($\tau_{\text{shock}} = 0.1$), the cumulative effects range from -0.5% for individuals initially at the 10th percentile to -1.4% for those initially at the 90th percentile. For positive shocks ($\tau_{\text{shock}} = 0.9$), the cumulative effects are larger for individuals starting at lower quantiles than for those at higher quantiles. These findings are consistent with the nonlinear persistence in η_{it} and underscore the importance of accounting for heterogeneous earnings dynamics when assessing the long-term consequences of economic shocks.

Figure 23: Impulse Response: Individual Labor Earnings

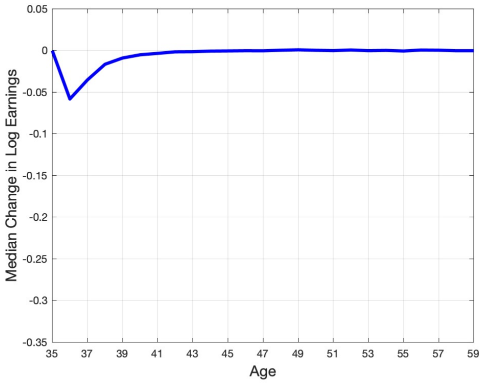

(a) $\tau_{init} = 0.1$, $\tau_{shock} = 0.1$

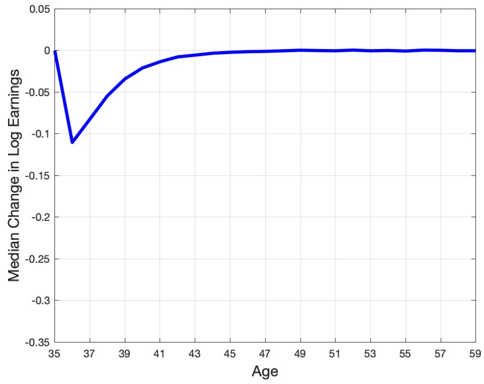

(b) $\tau_{init} = 0.9$, $\tau_{shock} = 0.1$

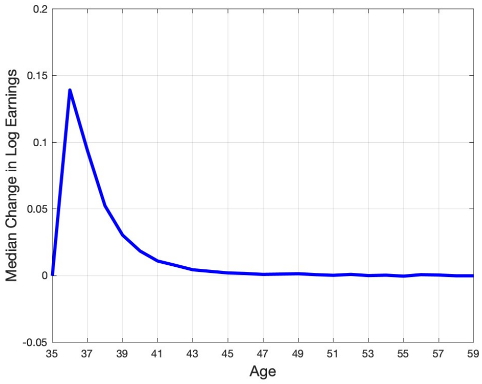

(c) $\tau_{init} = 0.1$, $\tau_{shock} = 0.9$

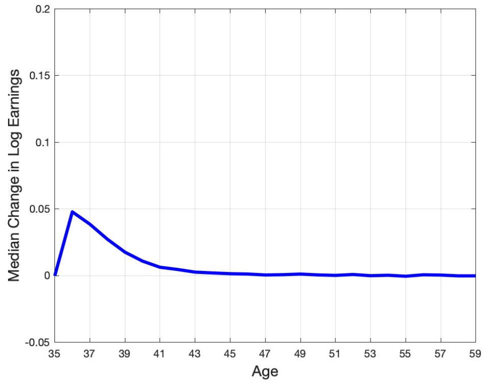

(d) $\tau_{init} = 0.9$, $\tau_{shock} = 0.9$

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 24: Cumulative Effects of Persistent Shocks: Individual Labor Earnings

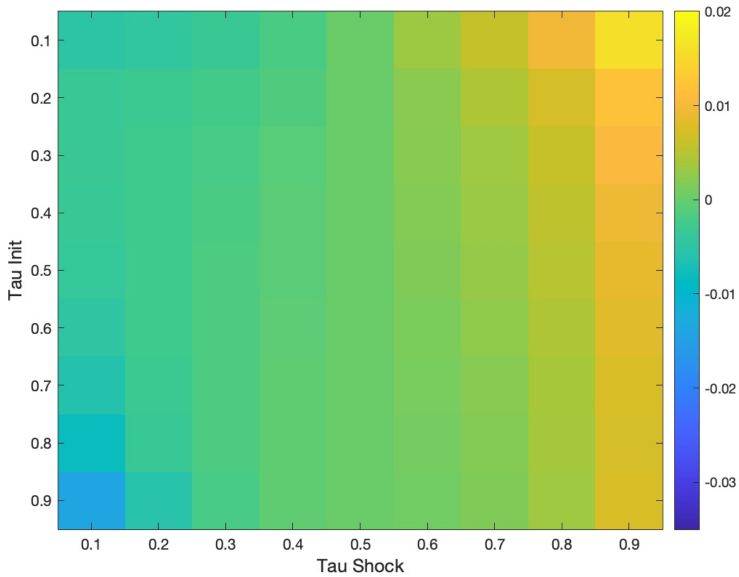

Notes. Individual-level data for 2011-2021 for a municipality.

6.2 Impulse Response Functions Using Household-level Data

Figure 25 presents the impulse response functions of persistent shocks for various household-level income measures, illustrating the dynamic impact of shocks that occur at age 36. A key empirical pattern is that the impact of a persistent shock is consistently larger for equivalized household earnings, total income, and after-tax income than it is for the head's earnings alone. This difference is particularly pronounced in extreme 'reverse' situations such as a large negative shock hitting a high-earning household (Panel b) or a large positive shock hitting a low-earning household (Panel c)where the conditional persistence of past income is at its lowest. The responses across the three different household-level aggregates are quantitatively similar to one another. Reflecting these dynamic responses, the long-term cumulative effects of shocks, shown in Figure 26, are also larger for equivalized household income than for the head's earnings. Furthermore, the impulse response and cumulative effects for the household head's labor earnings are found to be qualitatively and quantitatively similar to those for the broader individual male sample.

This finding reveals an important distinction between cross-sectional inequality and dynamic risk. While our cross-sectional analysis shows that household-level inequality is lower than individual-level inequality–suggesting effective income pooling within households–the impulse response analysis presents a contrasting picture. The same percentile transition (e.g., from the 90th to 10th percentile of the innovation distribution) generates larger earnings changes for household earnings than for household-head earnings. This indicates that the household earnings distribution exhibits different dynamic properties, with percentile movements representing larger economic magnitudes at the household level.

Several factors may explain this pattern. First, the persistent shocks captured by our model may not be purely idiosyncratic to one individual but rather reflect correlated shocks affecting multiple household members. When a persistent shock occurs-such as an industry- or occupation-specific downturn-it may simultaneously impact both spouses' earnings trajectories if they work in related sectors or share similar skill sets. Second, household responses to shocks may involve joint labor supply decisions that amplify the initial impact. For instance, when one spouse experiences job loss due to health problems or disability, the other spouse might reduce work hours to provide care and support. Similarly, when elderly parents require nursing care or a child develops special needs, both spouses may need to adjust their work schedules, leading to correlated earnings declines. These family care responsibilities create situations where an adverse shock to one household member necessitates labor supply reductions by other members. Third, demographic changes triggered by economic shocks such as fertility timing or household formation decisions affect the dynamics of equivalized income through changes in the equivalization factor, which we define as the square root of the household size.

It is important to note, however, that our data do not include information on public transfers. This omission may be significant, as it could lead to an overstatement of the volatility of cash inflows. For example, events like childbirth may cause a drop in parental earnings as they take leave or reduce their work hours, but the associated income loss would be mitigated by public support like childcare allowances. By not observing these transfers, our measure of equivalized household income may appear more volatile than it truly is, potentially overstating the magnitude of household-level responses to shocks.

Figure 25: Impulse Response: Household Earnings and Income

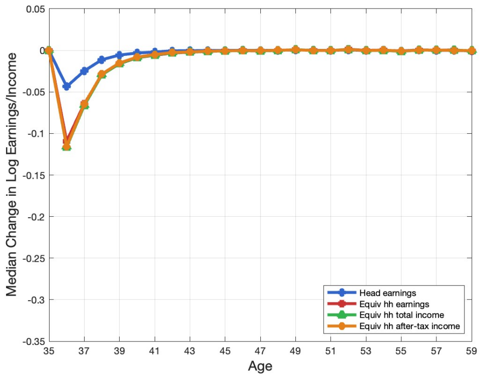

(a) $\tau_{init} = 0.1$, $\tau_{shock} = 0.1$

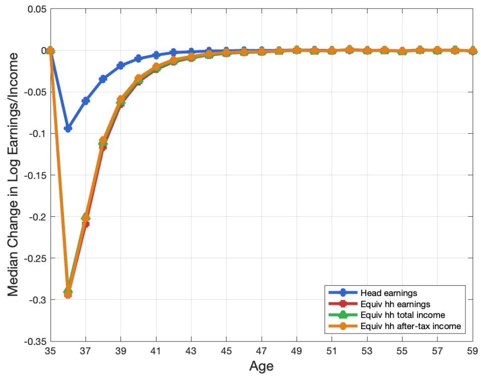

(b) $\tau_{init} = 0.9$, $\tau_{shock} = 0.1$

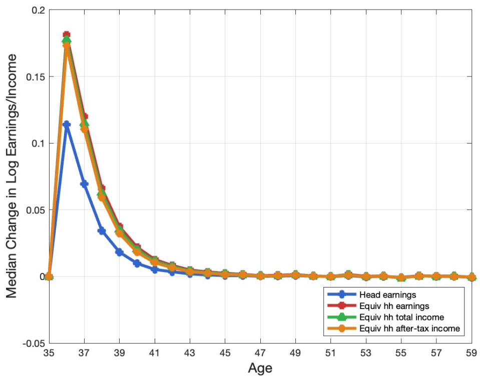

(c) $\tau_{init} = 0.1$, $\tau_{shock} = 0.9$

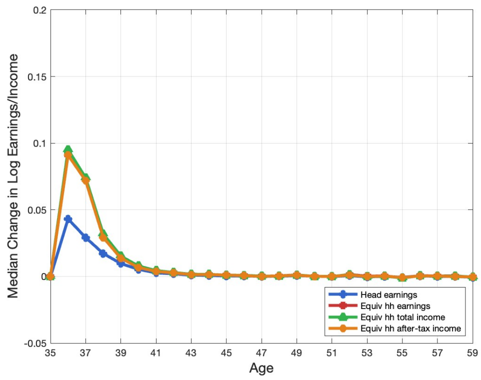

(d) $\tau_{init} = 0.9$, $\tau_{shock} = 0.9$

Notes. Household-level data for 2011-2021 for a municipality.

7 Concluding Remarks

This paper has utilized a unique administrative panel dataset of local tax records from a Japanese municipality to provide a detailed analysis of life-cycle earnings dynamics and inequality. By applying a flexible, nonlinear quantile regression framework, we offer several key contributions to the understanding of idiosyncratic risk in Japan.

First, we document significant nonlinearities in the earnings process that cannot be captured by standard linear models. The persistence of earnings is state-dependent, drop-

Figure 26: Cumulative Effects of Persistent Shocks: Household Earnings and Income

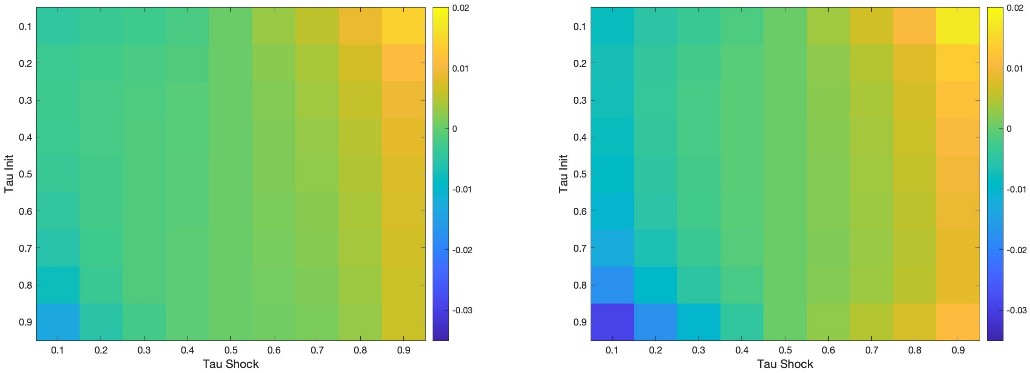

(a) Household head earnings (b) Equivalized household earnings

Notes. Household-level data for 2011-2021 for a municipality.

ping sharply for individuals who experience large 'reversal' shocks, such as a high-earner facing a major negative shock. This finding, consistent with evidence from the United States and Norway, suggests that events like career transitions or job displacements can fundamentally reset an individual's earnings trajectory, a crucial feature for accurately modeling lifetime income risk.

Second, our analysis reveals a nuanced distinction between cross-sectional risk sharing and dynamic insurance within households. Cross-sectional evidence confirms that households effectively pool income sources: inequality is significantly lower for equivalized household income than for individual earnings, and its J-shaped life-cycle profile contrasts sharply with the monotonically rising inequality observed for individuals. However, our dynamic analysis reveals that this pooling does not translate into insurance against persistent shocks. The same percentile innovations generate larger impulse responses for equivalized household income than for the head's earnings alone, indicating that the household income distribution has different dynamic propertieswith percentile transitions representing larger economic changes. This pattern likely reflects multiple mechanisms: correlated shocks affecting both spouses (such as industry-specific downturns), joint labor supply responses to adverse events (such as care responsibilities arising from health shocks), and demographic adjustments that affect household composition.

These findings carry important implications for economic modeling and policy. The evidence of nonlinear persistence underscores the need for models that move beyond linear-Gaussian assumptions to accurately assess welfare and design optimal social insurance systems. Furthermore, the dichotomy between cross-sectional smoothing and dynamic amplification at the household level highlights the necessity of looking beyond static measures to understand the true nature of intra-household insurance mechanisms.

Our analysis, however, is subject to several limitations that suggest avenues for future research. First, our data are from a single municipality and, while comprehensive for that locality, are not nationally representative. Replicating this analysis with data from other municipalities or, should it become available, national administrative data would be essential to confirm the generality of our findings. Second, our administrative records lack information on public transfers. This omission is particularly relevant for the household-level analysis, as unobserved transferssuch as employment insurance benefits following a job loss or childcare allowancescould mitigate the income volatility we measure, potentially tempering the shock amplification we observe. Finally, our dataset does not include information on consumption or wealth. Integrating a rich earnings process like the one estimated here into a structural life-cycle model with consumption and savings decisions would allow for a direct quantification of the welfare costs of nonlinear earnings risk and a more complete understanding of households' behavioral responses.

Appendix A Data Construction

Our analysis is based on local tax data derived from administrative tax records maintained by municipalities participating in the Local Government Data Development Project coordinated by the CREPE at the University of Tokyo. These data include detailed, individual-level information required to compute local inhabitant tax liabilities, encompassing various sources of income, statutory deductions, and basic demographic characteristics such as age and gender.

A key strength of the dataset is its administrative origin, which ensures comprehensive coverage of the resident tax population within each municipality, including individuals whose income is subject to the withholding tax system as well as those who file a final tax return. Each record contains a unique individual identifier and a municipality code, enabling the construction of a panel dataset that tracks individuals across multiple years within the same municipality.

A particularly distinctive feature of the dataset, rare among administrative tax records, is the inclusion of a household identifier and each individual's relationship to the household head. This structure allows us to construct a household-level panel, identify household heads, spouses, and dependents, and examine intra-household dynamics over time. This household linkage is critical for analyzing income aggregation, tax incidence, and the distributional impact of tax policies at the household level.

A.1 After-tax Income Measure

This section outlines the construction of after-tax income using our local tax data. Our calculation accounts for the complex structure of the Japanese tax system, including national income tax, local inhabitant tax, social insurance contributions, and various income deductions. Our after-tax income measure is constructed as follows:

$$\begin{aligned}\text{After-tax income} = {} & \text{Total income} + \text{Employment income deduction} \\ & - \text{National income tax} - \text{Local income tax} \\ & - \text{Social insurance contributions}\end{aligned} \tag{3}$$

This approach treats the employment income deduction as an income deduction for tax purposes rather than as estimated work expenses, which is consistent with how the Japanese tax system conceptualizes this deduction. Our calculator truncates taxable income to the nearest 1,000 yen as required by Japanese tax law and incorporates all major tax reforms during the study period (2011-2025), including changes to tax brackets, rates, and deduction schedules.

Total Income. We define total income as the sum of all taxable income sources subject to comprehensive taxation, excluding interest and dividend income. We exclude interest and dividend income because they are typically taxed through withholding or separate self-assessed taxation. The "net" designation indicates that necessary expenses and statutory deductions such as the employment income deduction have been deducted.

Employment Income Deduction (EID). We apply the statutory employment income deduction schedule, which varies by employment income (*kyuyo shunyu*) level and year. The deduction ranges from 550,000-650,000 yen for low-income earners to caps of 1.95-2.45 million yen for high-income earners, depending on the tax year.

National Income Tax. We calculate taxable income by subtracting all applicable deductions from total net income. Our local tax data include income deductions such as the basic deduction, exemption for dependents, spousal deduction, special spousal deduction, casualty loss deduction, medical expense deduction, and social insurance premium deduction, among others. Although these deduction amounts are defined for local tax purposes and may slightly differ from those for national income tax, we apply the same values to both national and local tax calculations, except for the basic deduction. Because these differences are small, we believe that this simplification does not significantly affect our results. We recalculate the basic deduction for national income tax based on its local tax counterpart, taking advantage of the fact that both systems share the same income brackets for determining eligibility. This adjustment is particularly important because the basic deduction, together with the employment income deduction, determines the tax-exempt income threshold.

National income tax liability is then computed using the Quick Calculation Table of Income Tax, which reflects Japans progressive tax schedule with marginal rates ranging from 5% to 45% (post-2015), along with corresponding deduction amounts. Since 2013, we include the 2.1% Special Income Tax for Reconstruction.

Local Inhabitant Tax. We implement a simplified model with a flat 10% rate on local taxable income plus a per-capita levy of 5,000 yen. This approach captures the essential features of the local tax system while abstracting from minor inter-municipal variations.

Limitations. The main limitation of our after-tax income calculation is the lack of information on tax credits under the national income tax system and on public transfers such as social security benefits. These omitted components, particularly public transfers, are likely to play an important role in income redistribution and may lead to an understatement of the redistributive effects in our analysis.

A.2 Household-level Data

This section explains how raw individual-level local tax data are aggregated into the household-year panel used in the empirical analysis. Households are uniquely identified by combining household identifiers with municipality codes. This creates a consistent household identifier that allows us to track households across years.

For each household, we identify the household head using the following sequential criteria: (1) Individual designated as household head (*setai-nushi*) in the records and currently living in the household; (2) Highest income among candidates[3]; (3) Male gender (in case of income ties); (4) Oldest age (in case of remaining ties). Households without a living head are excluded from the analysis.

We then aggregate individual-level data to the household level, creating a household-year panel. First, we calculate the sum of monetary variables for all household members, including pretax earnings, total net income, and after-tax income. Second, we construct household composition variables, including the number of household members, the number of children, the number of adults, the number of income earners, the presence of a cohabiting spouse, etc. To account for household size and composition, we create equivalized income measures by dividing household income by the square root of household size.

We classify households into six distinct types based on marital status, presence of children, and employment status: (1) Single without children; (2) Single with children (single parents); (3) Married couples with both spouses working, without children; (4) Married couples with both spouses working, with children; (5) Married couples with one working spouse, without children; (6) Married couples with one working spouse, with children. Working status is determined based on positive income, while spousal deduction receipt is used to identify households where one spouse has limited income. Head- and spouse-specific information (age, sex, own income, spousal deductions) is merged back into the household record.

A.3 Sample Selection

In the benchmark analysis, we focus on a single municipality in Japan, covering the period from 2011 to 2021. Our individual-level analysis employs several sequential sample restrictions. Following Guvenen et al. (2021), we select individuals who meet the following criteria: 1. residing in the municipality on January 1 of each year; 2. aged 25-59; 3. male; 4. labor earnings above the minimum threshold, defined as one quarter of full-time work (13 weeks at 40 hours per week) at the lowest minimum wage across all prefectures; 5. primary source of income is employment income. We identify these individuals by

[3] For the municipalities studied in this paper, criterion (2) is sufficient to assign a unique household head to each household.

comparing the mean of different income sources (business, farm, real estate, interest, dividend, and other earnings) across all observed years and selecting those for whom salary income is both the maximum income source and positive. This approach ensures that we focus on individuals with stable attachment to the labor market as employees rather than self-employed individuals or those with primarily non-labor income.

For the household-level analysis, we restrict our sample to households that maintain the same household head throughout the panel period. Because our local tax dataset tracks only individuals residing in the municipality, we observe changes in household head when the head moves out of the municipality due to his/her job assignment. To minimize the impact of this issue, we focus on households with stable heads. Our final household-level sample consists of households with stable male heads aged 25–59, whose labor earnings exceed the minimum threshold defined above and whose primary income source is employment income.

Appendix B Additional Figures

B.1 Alternative Municipality

Table 4: Descriptive Statistics of the Cross-Section Sample

	Obs	Mean	Std. Dev.	P10	P50	P90
Individual-level sample						
Age	1129101	42.30	9.65	29	43	56
Labor earnings	1129101	5439.85	3530.19	2238.35	4830.71	8797.73
Total income	1129101	5468.41	3651.18	2246.09	4838.08	8827.59
After-tax income	1129101	4194.50	2369.46	1846.32	3822.96	6627.11
Household-level sample						
Age	930352	43.19	9.47	30	44	56
Household size	930352	2.70	1.43	1.00	3.00	4.00
Labor earnings	930352	7127.68	4618.30	2848.38	6299.21	12044.90
Total income	930352	7182.29	4772.76	2863.80	6325.56	12113.60
After-tax income	930352	5562.23	3252.78	2333.80	5032.82	9272.09
Labor earnings (equivalized)	930352	4551.21	2705.30	2126.63	4069.93	7323.36
Total income (equivalized)	930352	4582.53	2791.14	2141.07	4085.13	7358.30
After-tax income (equivalized)	930352	3549.57	1863.67	1755.37	3238.50	5590.80

Notes. Columns designated as Obs, Mean, Std. Dev., P10, P50, and P90 report the number of observations, mean, standard deviation, and the 10th, 50th (median), and 90th percentiles of the respective variable. The monetary values are deflated by the Core CPI (excluding fresh food, 2020 base year) and expressed in thousands of yen. For the household-level sample, the equivalized income is calculated by dividing the household income by the square root of household size.

Figure 27: Inequality Measures

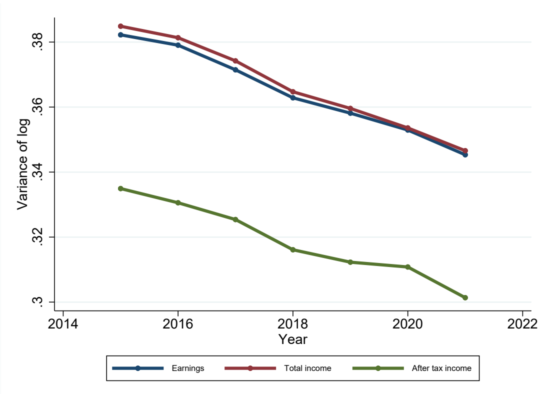

(a) Variance of log

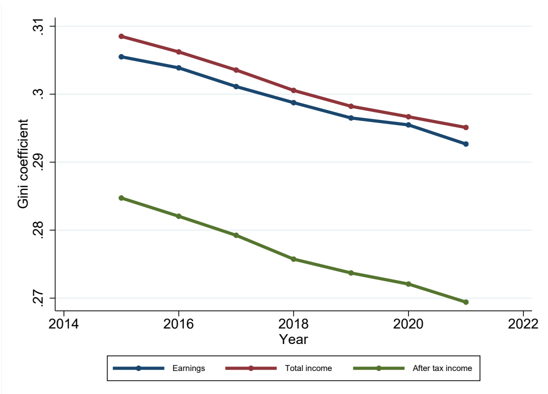

(b) Gini

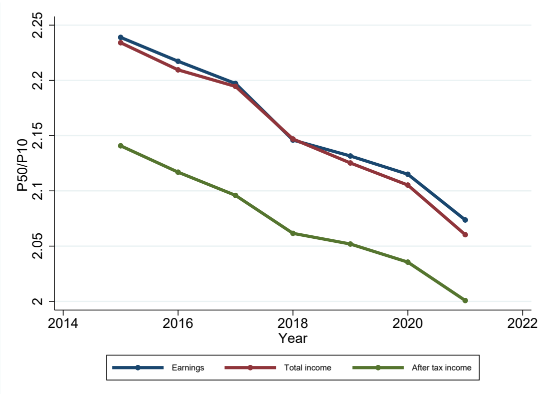

(c) P50/P10

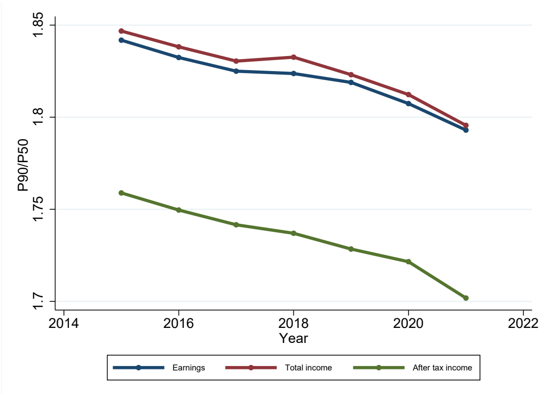

(d) P90/P50

Notes. Individual-level data are used.

Figure 28: Inequality Measures

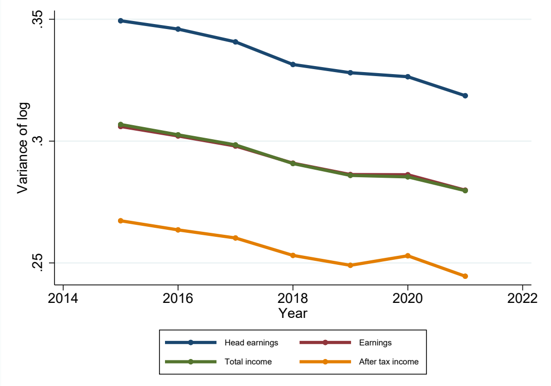

(a) Variance of log

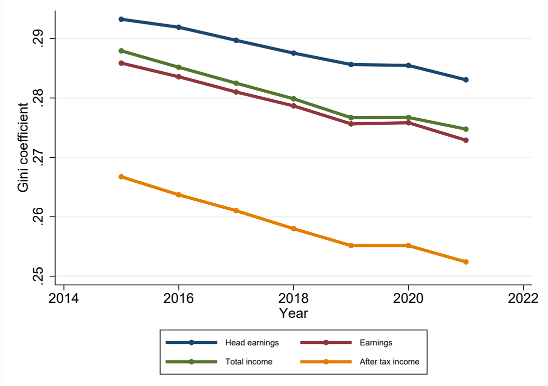

(b) Gini

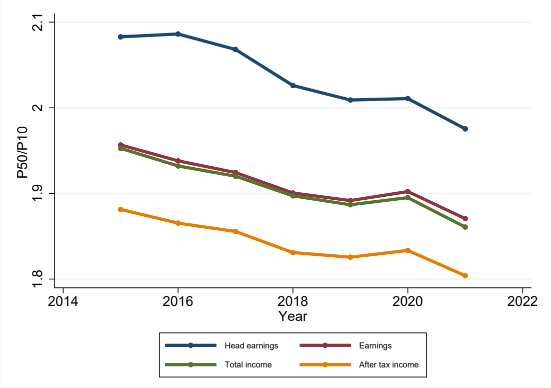

(c) P50/P10

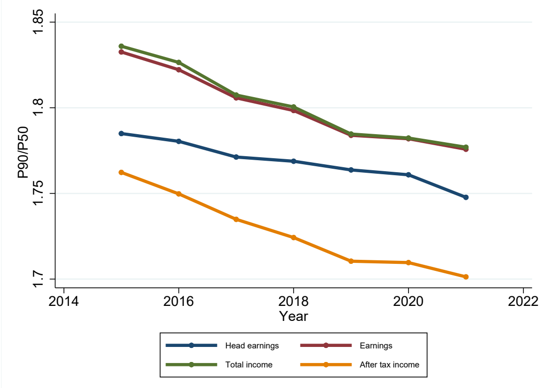

(d) P90/P50

Notes. Household-level data are used.

Figure 29: Role of Income Deductions

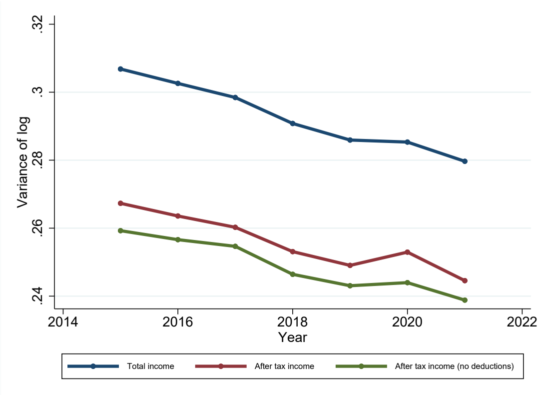

(a) Variance of log

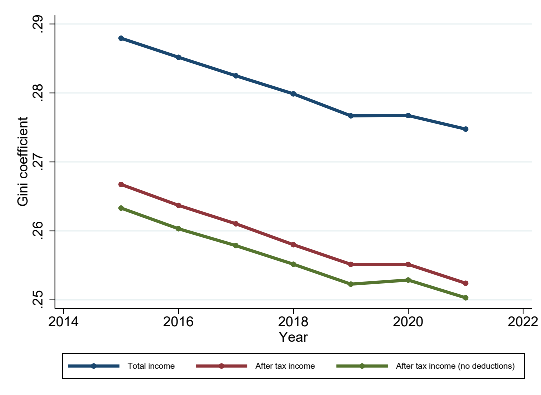

(b) Gini

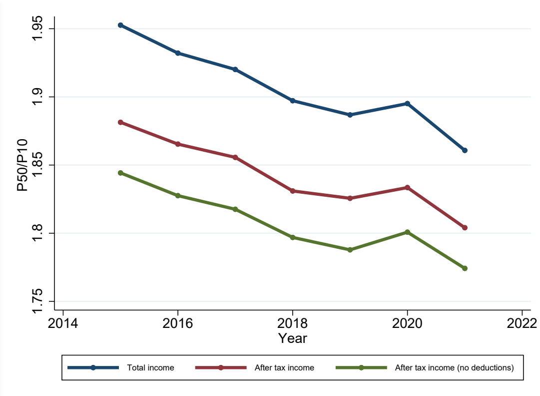

(c) P50/P10

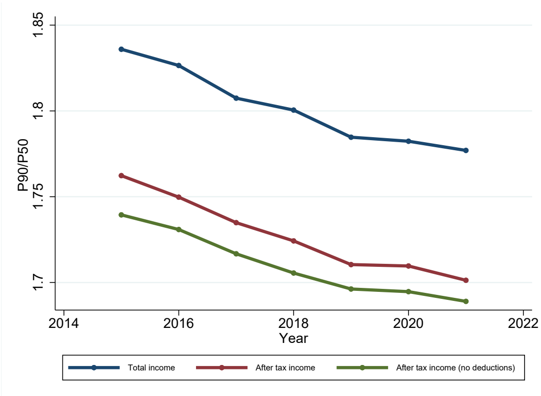

(d) P90/P50

Notes. Household-level data are used.

Figure 30: Lifecycle Inequality

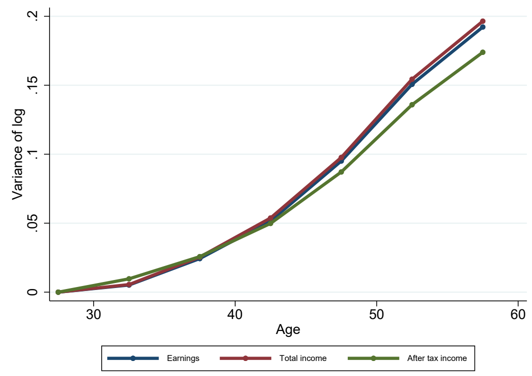

(a) Time effects, Individual

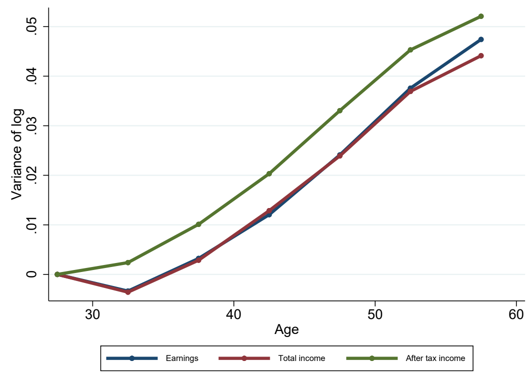

(b) Cohort Effects, Individual

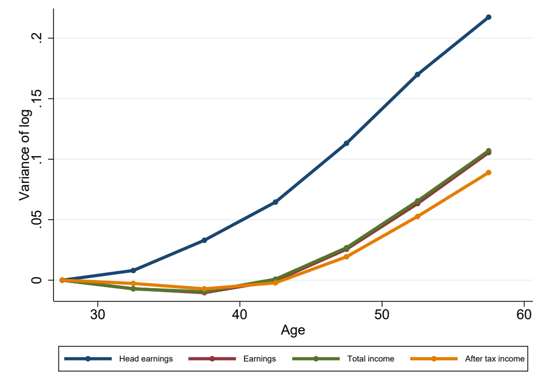

(c) Time Effects, Household

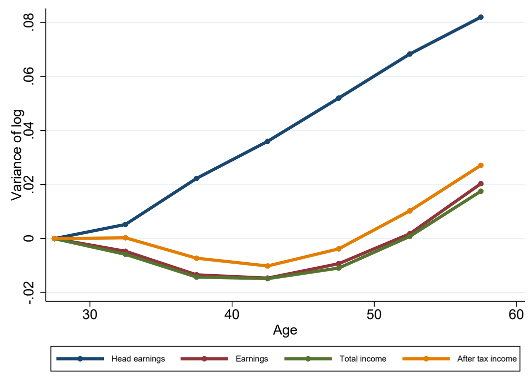

(d) Cohort Effects, Household

Notes. Panels (a) and (b) show individual-level data, while (c) and (d) show household-level equivalized data.

Figure 31: Persistence of Log Earnings Residuals

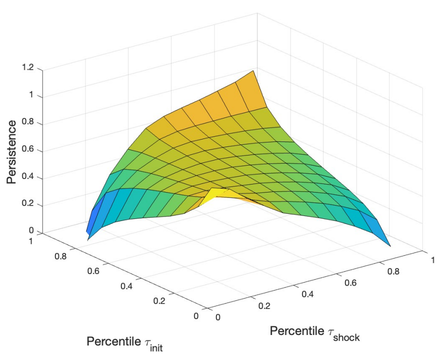

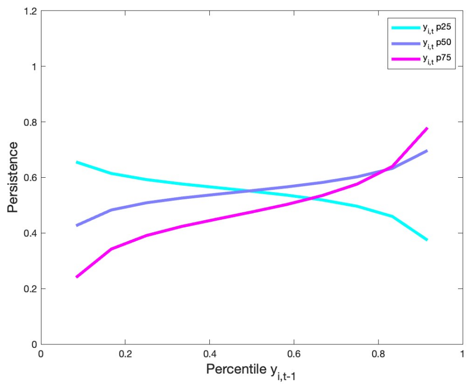

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 32: Persistence of Log Earnings Residuals By Age

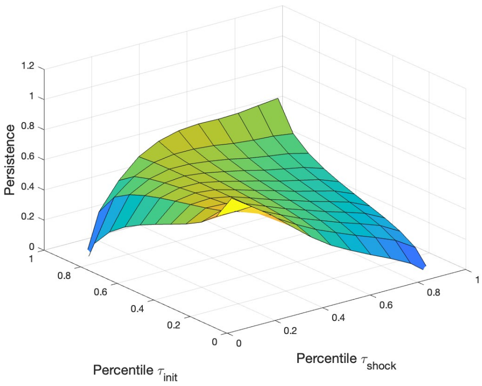

(a) 25 years old

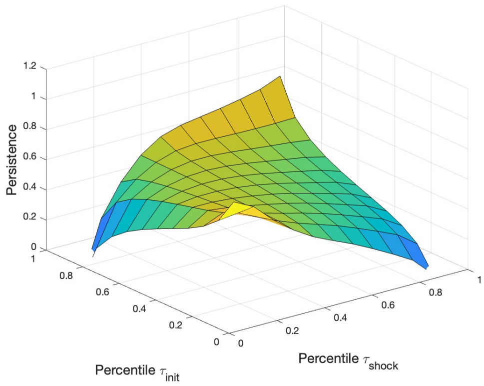

(b) 35 years old

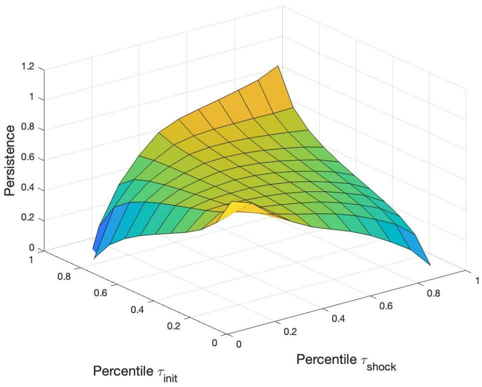

(c) 45 years old

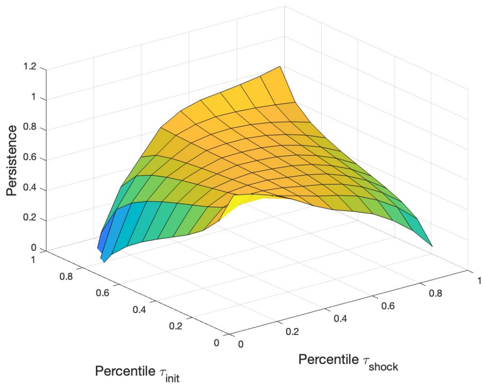

(d) 55 years old

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 33: Persistence of Log Earnings Residuals

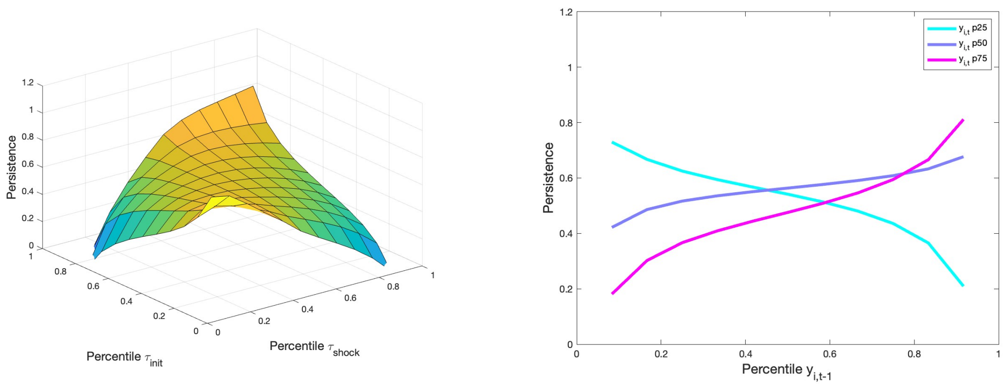

Notes. Household-level data for 2011-2021 for a municipality.

Figure 34: Persistence of Log Earnings Residuals By Age

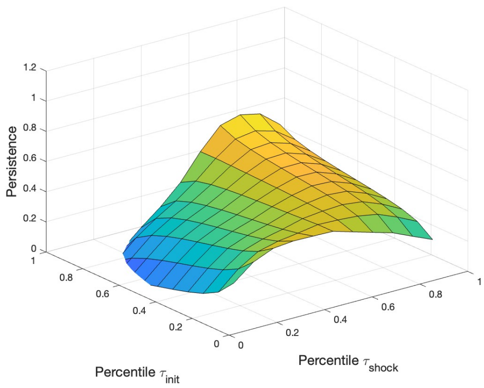

(a) 25 years old

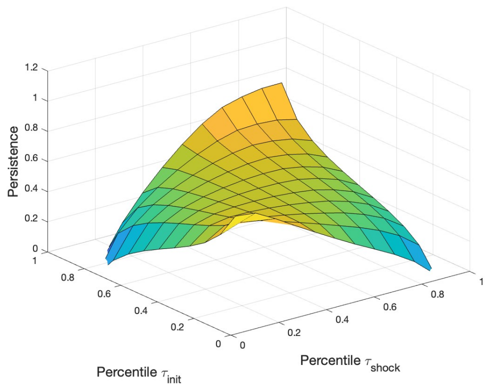

(b) 35 years old

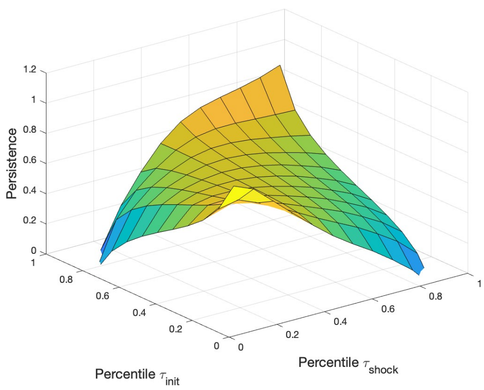

(c) 45 years old

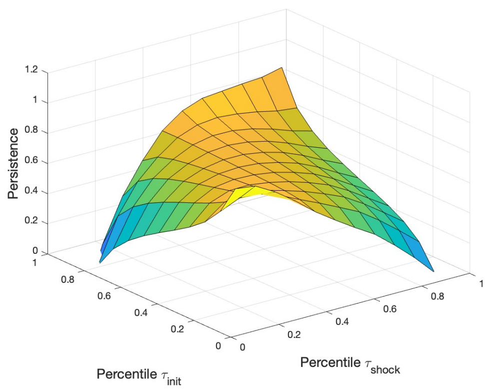

(d) 55 years old

Notes. Household-level data for 2011-2021 for a municipality.

Figure 35: Conditional Dispersion of η_{it}

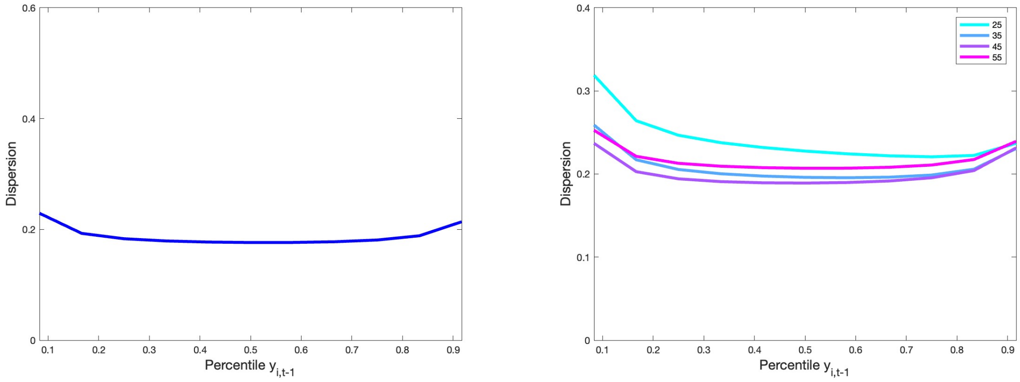

(a) Whole Sample (b) Age Group

Notes. Individual-level data for 2011-2021 for a municipality.

B.2 Estimation of The Nonlinear Model (2)

Table 5: Descriptive Statistics of the Estimation Sample

	Obs	Mean	Std. Dev.	Min	Max
Individual-level data					
Age	60,269	43.85	6.89	25	59
Labor earnings (Residual)	60,269	0.00	0.17	-2.01	1.08
Total income (Residual)	60,269	0.00	0.17	-1.86	1.17
After-tax income (Residual)	60,269	0.00	0.16	-2.11	1.01
Household-level data					
Household head's age	51,645	44.36	6.66	25	59
Labor earnings (Residual)	51,645	0.00	0.17	-1.73	0.89
Total income (Residual)	51,645	0.00	0.17	-1.87	0.89
After-tax income (Residual)	51,645	0.00	0.16	-1.92	0.83

Figure 36: Conditional Skewness of η_{it}

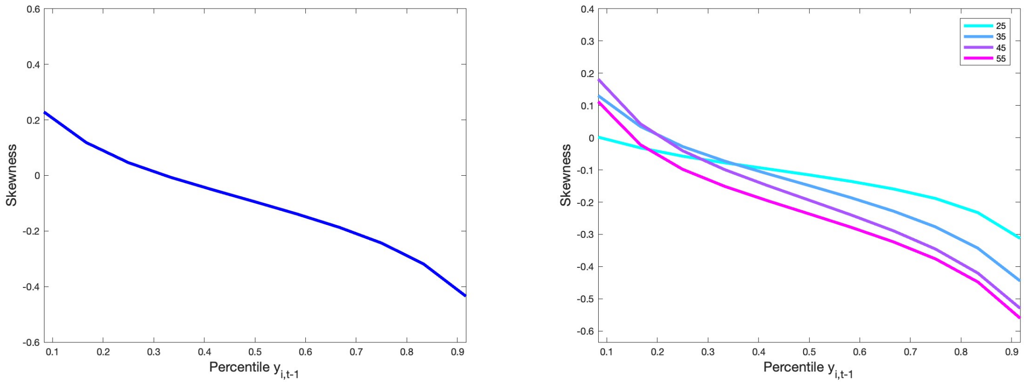

(a) Whole Sample (b) Age Group

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 37: Conditional Kurtosis of η_{it}

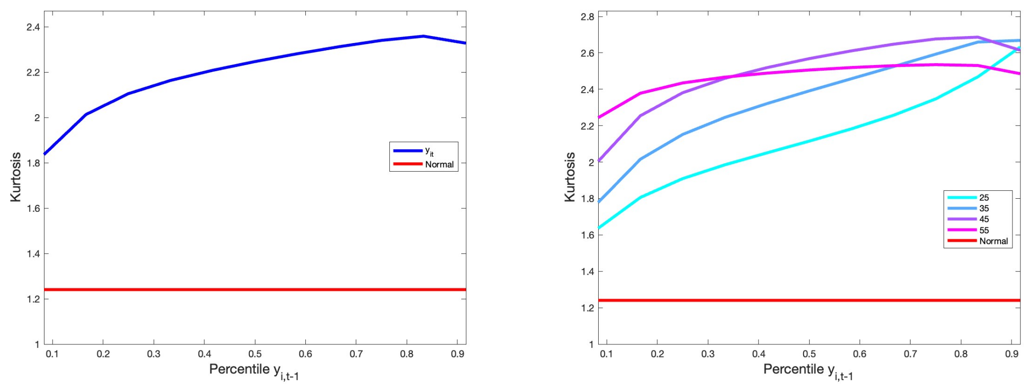

(a) Whole Sample (b) Age Group

Notes. Individual-level data for 2011-2021 for a municipality.

Figure 38: Conditional Dispersion of ε_{it}

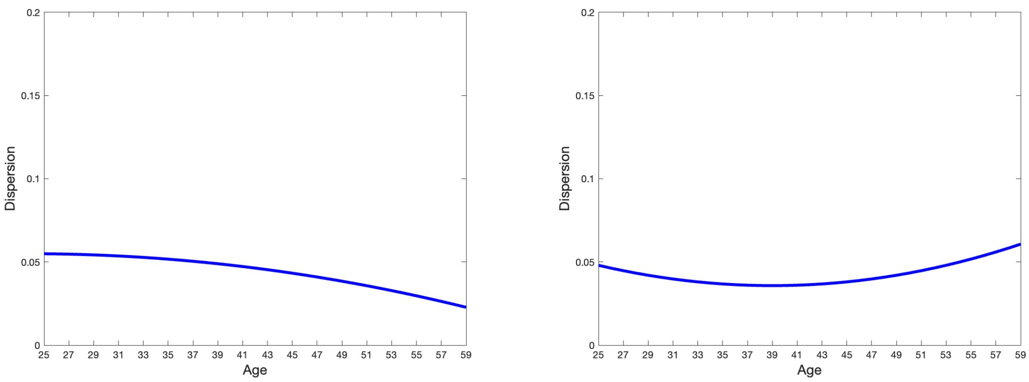

(a) Individual (b) Household

Figure 39: Conditional Skewness of ε_{it}

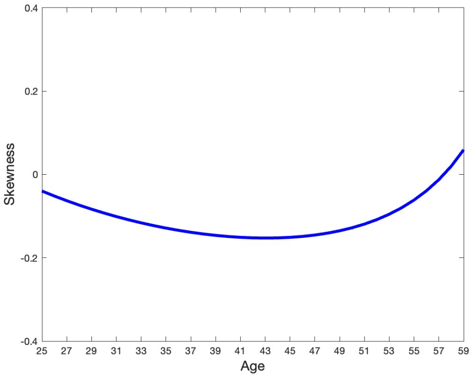

(a) Individual

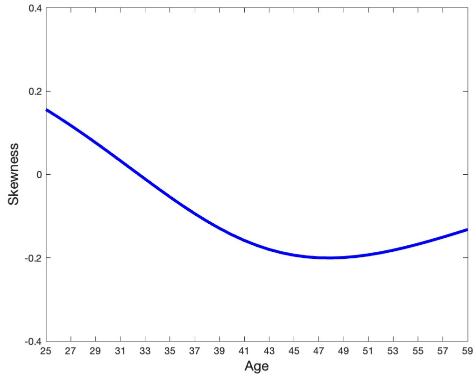

(b) Household

Figure 40: Conditional Kurtosis of ε_{it}

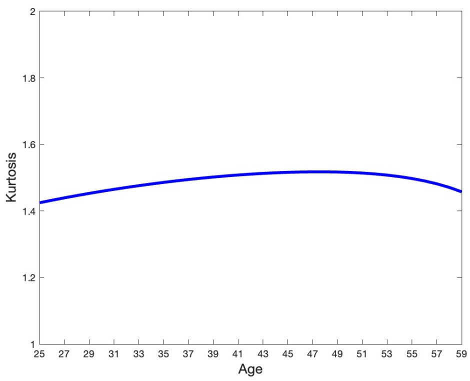

(a) Individual

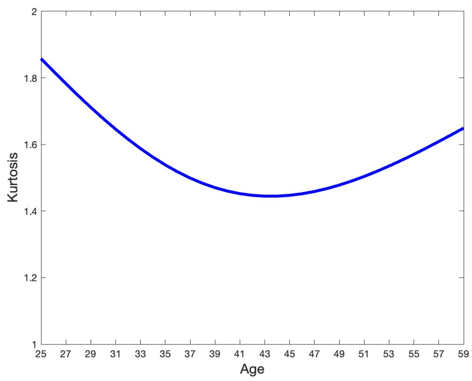

(b) Household

References

Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica 85*(3), 693–734.

Arellano, M., R. Blundell, S. Bonhomme, and J. Light (2024). Heterogeneity of consumption responses to income shocks in the presence of nonlinear persistence. *Journal of Econometrics 240*(2), 105449.

Arellano, M. and S. Bonhomme (2016). Nonlinear panel data estimation via quantile regressions. *Econometrics Journal 19*(3), C61–C94.

Blundell, R., M. Graber, and M. Mogstad (2015). Labor income dynamics and the insurance from taxes, transfers, and the family. *Journal of Public Economics 127*, 58–73.

Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010). Quantile and probability curves without crossing. *Econometrica 78*(3), 1093–1125.

De Nardi, M., G. Fella, and G. Paz-Pardo (2020). Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association 18*(2), 890–926.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What do data on millions of U.S. workers reveal about life-cycle earnings? *Econometrica 89*(5), 2303–2339.

Guvenen, F., L. Pistaferri, and G. L. Violante (2022). Global trends in income inequality and income dynamics: New insights from GRID. *Quantitative Economics 13*(4), 1321–1360.

Hu, Y. and S. M. Schennach (2008). Instrumental variable treatment of nonclassical measurement error models. *Econometrica 76*(1), 195–216.

Wilhelm, D. (2015). Identification and estimation of nonparametric panel data regressions with measurement error. Cemmap Working Paper CWP34/15.