



Network for Studies on Pensions, Aging and Retirement

Household Earnings Risk and its Impact on Consumption and Portfolio Decisions

Gonzalo Paz Pardo

NETSPAR ACADEMIC SERIES

PhD 10/2020-003

Household Earnings Risk and its Impact on
Consumption and Portfolio Decisions

Gonzalo Paz Pardo

University College London (UCL)

Submitted towards the degree of PhD in Economics

I, Gonzalo Paz Pardo, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

This thesis includes elements from my papers [De Nardi, Fella, Knoef, Paz-Pardo and Van Ooijen \(2018\)](#), [De Nardi, Fella and Paz-Pardo \(2019\)](#), and [Paz-Pardo \(2019\)](#). While Chapter 3 is fully my own work, Chapters 1 and 2 are derived from joint research.

In Chapter 1, I contributed by proposing the measures to study and the empirical approach, and writing the code that computes all of the measures in the Dutch administrative data, which was cleaned and organised by my coauthors.

In Chapter 2, I contributed by performing the data analysis, estimating the earnings process, building and calibrating the structural model, and providing a first analysis of the results. My coauthors and I jointly defined the research questions and developed the discretization method we propose.

Abstract

Household labor earnings are unequal and risky. Their persistence and the distribution of the shocks they face depend on age, birth cohort, and the position of a household in the earnings distribution. Understanding these features is key to explain household consumption decisions, the allocation of household savings between different asset classes, and the role of different insurance mechanisms in smoothing out those risks. Using US survey data and Dutch administrative data, I document that household earnings are less persistent for the young and for the income-poor and that shocks to earnings are infrequent and negatively skewed. Although the tax and transfer system insures part of these risks, particularly in the Netherlands, they are also present in disposable income. I then turn to evaluating the implications of these rich features by comparing, in the context of a standard life-cycle model, a flexible earnings process that incorporates them against a canonical process with constant persistence and normal shocks. I find that considering richer earnings dynamics helps us to better understand the evolution of cross-sectional consumption inequality over the life-cycle and the pass-through of persistent earnings shocks to consumption. Many of these features of earnings have changed over time. Using US data, I document that earnings are more unequal and riskier for younger generations. I argue that lower initial and lifetime earnings for the income-poor and larger earnings variability across the board are key to explain the reduction in home-ownership rates between the generations born in the 1940s and 1980s. I show the relevance of this mechanism in a flexible life-cycle model with risk-free assets, stocks, houses, and mortgages, and correlated idiosyncratic and aggregate risks. Changes in financial constraints also matter: looser mortgage requirements helped the young buy houses, and lower participation costs rationalize higher stock market participation.

Impact statement

Understanding the risks faced by households and how they shape their decisions is key for the design of many relevant economic policies. Thus, the research contained in this thesis has potential impact both for economic research and for public policy design.

Looking at the impact within academia first, this thesis provides contributions from the empirical, methodological, and conceptual points of view. The facts documented in this thesis, regarding household earnings dynamics and portfolio compositions, both from an international and intergenerational perspective, can spur further research on their determinants and implications. The methods proposed in this thesis, both in terms of richer earnings dynamics processes and richer structural life-cycle models, can be the building blocks for further research on distributional issues. Additionally, the estimated processes can be used as inputs that other researchers may include in their models. Finally, this thesis advances and quantifies several important causation channels for the drop in homeownership rates and changes in household portfolios, which can inform future work on the topic.

Additionally, this thesis can inform public policy. Governments in many countries are worried about the decline in homeownership rates and the inability of younger households to buy houses. This project suggests that changes in their labor market outcomes have been key in explaining why they don't buy houses, so policies that insure workers against labor market income fluctuations and that correct earnings inequality are likely to have an impact on homeownership rates. Additionally, I study other policy interventions, such as the reduction in costs of access to financial markets, that can be useful for a policymaker who is concerned about increasing levels of wealth inequality.

More broadly, this thesis contributes to the burgeoning public discussion on income and wealth inequality, with a particular focus on intergenerational distribution. In particular, I show that the drop in homeownership rates for younger generations in the United States can be rationalized by changes in economic conditions, even if we assume that preferences with respect to owning a house have not changed over time.

Acknowledgements

A PhD is a lengthy endeavour. There are many people along the way that have contributed to this work, both directly and indirectly, to which I am indebted.

My deepest thanks go to my PhD supervisor, Mariacristina De Nardi, for all of her help with my work, but mostly for being a role model as a researcher, an advisor, and a mentor. I left all of our numerous meetings during the PhD with good suggestions, clearer ideas, and encouragement to continue my work. She is also responsible for all of the Oxford commas in this thesis.

I would also like to thank the several members of UCL faculty, students, and external researchers that took some of their time to provide very valuable feedback on this work.

I gratefully acknowledge financial support from the Economic and Social Research Council (ESRC) for my graduate studies, and financial support from Netspar for Chapter 1.

Last but not least, I'd like to thank my family and friends, both in the UK and in Spain, both those that sat through practice presentations and those helped me disconnect from work. A particular thought goes to those that never understood what my work and economic research is all about, but still supported me because they would believe in anything I devoted my time to. I hope I haven't disappointed them.

Contents

Abstract	5
Impact statement	7
Introduction	23
Chapter 1 Nonlinear Dynamics of Earnings in the Netherlands: Sources and Insurance Mechanisms	25
1.1 Introduction	25
1.2 Data and approach	27
1.3 Results	31
1.3.1 Male wages	31
1.3.2 Male earnings	34
1.3.3 Household insurance	39
1.3.4 Insurance from taxes and transfers	41
1.4 Conclusion	42
Chapter 2 The Implications of Nonlinear Earnings Dynamics for Household Self-Insurance and Welfare	47
2.1 Introduction	47
2.2 Earnings Data and their Features	51
2.2.1 Individual Pre-tax Earnings in the PSID and the W2 Data	53
2.2.2 Individual Pre-tax versus Household Disposable Earnings in the PSID	54
2.3 Earnings Processes and their Estimation	56
2.3.1 Earnings Processes	56
2.4 Estimating the Nonlinear Earnings Process	58
2.4.1 Estimating the Canonical Linear Earnings Process	59
2.4.2 Comparing the Implications of the Nonlinear and Canonical Earnings Processes	59
2.5 The Discretized Nonlinear Earnings Process	61
2.6 The Model	63
2.6.1 Demographics	63
2.6.2 Preferences and Endowments	63
2.6.3 Markets and the Government	64

2.6.4	The Household's Problem	64
2.6.5	The Model Calibration	65
2.7	Consumption, Wealth, and Welfare Implications	66
2.7.1	Consumption Inequality over the Working Life	66
2.7.2	Measuring Self-Insurance against Earnings Shocks	72
2.7.3	Wealth	76
2.7.4	Welfare	77
2.8	Conclusions	78
Chapter 3	Intergenerational Changes on Household Earn- ings Risk, Homeownership, and Household Port- folios	81
3.1	Introduction	81
3.2	An overview of intergenerational changes	88
3.2.1	Distribution of earnings	88
3.2.2	Earnings risk	89
3.2.3	Housing prices and stock returns	91
3.2.4	Financial conditions	91
3.2.5	Homeownership and portfolio composition	92
3.3	A business-cycle dependent earnings process	93
3.3.1	Implications of the earnings process	96
3.3.1.1	Intergenerational differences	96
3.3.1.2	State-dependence of the earnings process	97
3.4	Model	99
3.4.1	Demographics	100
3.4.2	Preferences	100
3.4.3	Environment and technologies	101
3.4.3.1	Aggregate state	101
3.4.3.2	Earnings	102
3.4.3.3	Liquid accounts	102
3.4.3.4	Risky financial assets	102
3.4.3.5	Housing	103
3.4.3.6	Mortgages	104
3.4.4	The government	105
3.4.5	Timing	106
3.4.6	Budget constraint	107
3.4.7	Household's problem	107

3.5	Calibration	109
3.5.1	Aggregate state	109
3.5.2	Externally calibrated parameters	112
3.5.3	Internally calibrated parameters, targets, and model fit	114
3.6	Results	116
3.6.1	Untargeted moments, 1940s cohort	116
3.6.2	Explaining intergenerational differences in homeownership	119
3.6.2.1	Decomposing the decrease in homeownership	120
3.6.3	Explaining the changes in stock market participation	125
3.6.4	Adjusting housing prices	126
3.7	Implications	128
3.7.1	Reaction to shocks: consumption and homebuying decisions	128
3.7.1.1	Generational changes	128
3.7.1.2	Shocks to the aggregate state and the Great Recession	129
3.7.2	The future of the 1980s generation	131
3.8	Extensions and alternative specifications	134
3.8.1	Canonical earnings process	134
3.8.2	Other extensions	135
3.9	Conclusion	136
A	Appendix to Chapter 1	155
1.1	Dutch male wages, computed using actual hours worked.	155
1.2	Non-robust measures of skewness and kurtosis	156
1.3	Male earnings by recent earnings	157
1.4	Household earnings and pre-tax income	157
B	Appendix to Chapter 2	159
2.1	PSID data	159
2.1.1	The PSID	159
2.1.2	Sample Selection	159
2.1.3	Income Definition	160
2.1.3.1	Nonfinancial Pre-Tax Income	160
2.1.3.2	Tax Function	161
2.1.3.3	Equivalization	161
2.1.4	Estimating the Canonical and Nonlinear Earnings Processes	161
2.1.4.1	Estimating the Nonlinear Earnings Process	161

2.1.4.2	Estimating the Canonical Earnings Process . . .	162
2.2	Computation of the Variances of Log Earnings and Log Consumption.	162
2.3	Robust Measures of Skewness and Kurtosis	164
C	Appendix to Chapter 3	166
3.1	Data	166
3.1.1	PSID	166
3.1.1.1	Measures and sample selection	167
3.1.1.2	Tax progressivity	169
3.1.2	SCF	170
3.1.2.1	Bequest targets	171
3.1.3	Higher order moments of earnings	171
3.2	Computational appendix	172
3.2.1	EGM method	172
3.2.2	Global minimization	173
3.3	Intergenerational differences, alternative definitions	173
3.3.1	Sample composition	173
3.3.2	Choice of deflator	176
3.3.3	Marital dynamics and family composition	177
3.3.4	Earnings process, fit of variances	178
3.3.5	Earnings process, additional Figures	179
3.3.6	Portfolio composition, additional Figures	179
3.3.7	Equivalence scales by generation	180
3.4	Model, additional results and descriptions	180
3.4.1	Adjusting housing prices	180
3.4.2	Household reactions to aggregate shocks	183
3.4.3	Canonical process	184
3.5	Additional robustness checks	185
3.5.1	Initial zero wealth	185
3.5.2	Empirical initial wealth and inter-vivos transfers	187
3.5.3	Equilibrium effects on stock returns	188
3.5.4	Marital dynamics and family sizes	189
3.5.5	Per-period participation costs	190
3.5.6	Local correlation of income shocks and house prices	190
3.5.7	House sizes	192

List of Tables

1	Estimates for the canonical earnings process (standard errors in parentheses)	59
2	Insurance coefficients	74
3	Wealth distribution	77
4	Consumption measure of welfare costs	78
5	Externally calibrated parameters and sources	114
6	Targeted moments, model fit, and calibration	115
7	Contribution of each factor in the change in homeownership with respect to the 1940s generation (% of the change), by age	121
8	Wealth Gini at retirement, data vs model. The data is obtained from people aged 55-64 in the SCF (2001 and 2016 SCF for the 1940s and 1960s cohort respectively). For the model, standard errors from simulation are in parentheses.	134
B.1	Estimates for the canonical earnings process: cohort vs. year effects	164
C.1	Parameter estimates, business-cycle varying canonical process . .	184
C.2	Contribution of each factor in the change in homeownership with respect to the 1940s generation (% of the change), by age. Model with zero initial wealth.	186

List of Figures

1	Dutch male wages. Wage persistence (top left) and moments of wage changes: standard deviation (top right), skewness (middle left), Kelley’s skewness (middle right), kurtosis (bottom left), and Crow-Siddiqui kurtosis (middle right), by age group and previous earnings percentile.	32
2	NL, male wages (left) and male earnings (right). Persistence (top row), standard deviation (second row), Kelley’s skewness (third row), skewness (fourth row), and Crow-Siddiqui kurtosis (bottom row).	35
3	Variance of changes in male earnings, wages and hours. Top: all workers; bottom left: workers with a positive hours change; bottom right: workers with a negative hours change.	38
4	Skewness and kurtosis of changes in male earnings, wages and hours	38
5	Male earnings changes versus hours and wage changes. Each dot represents a decile of changes in male earnings. First decile of previous earnings (left), median decile of previous earnings (middle), 9th decile of previous earnings (right).	39
6	NL, household earnings (left) and disposable income (right). Persistence (top row), standard deviation (second row), Kelley’s skewness (third row), skewness (fourth row), and Crow-Siddiqui kurtosis (bottom row).	40
7	Male earnings changes and female labor supply. Each dot represents a decile of changes in male earnings. Lowest decile of previous male earnings (left), median decile of previous male earnings (center), 9th decile of previous male earnings (right)	41
8	NL, household earnings (left) and disposable income (right). Persistence (top row), standard deviation (second row), Kelley’s skewness (third row), skewness (fourth row), and Crow-Siddiqui kurtosis (bottom row).	43
9	Household before- vs after-tax income. Each dot represents a decile of changes in male earnings.	44

10	NL, age 55-59, Relative contribution of transfers and taxes to the standard deviation of household income. Red line, household gross income, gold line: including disability insurance, green line: also including unemployment insurance, dotted green line: also including social assistance, dotted blue line also including pensions, dotted red line: also net of taxes.	45
11	Standard deviation, skewness, and kurtosis of male pre-tax earnings growth in the W2 (top row) and PSID (middle row), and of household after-tax earnings growth in the PSID (bottom row) .	53
12	Persistence in log-earnings as a function of previous earnings rank and the rank of the shock received in the current period. PSID data.	55
13	Age dependence of second moments: nonlinear vs canonical process. Top left, standard deviation of the innovation to the persistent component. Top right, standard deviation of the transitory shock. Bottom left, autocorrelation of the persistent component. Bottom right, cross-sectional variance of log earnings. The dotted bands represent bootstrapped 95% confidence intervals.	60
14	Skewness and kurtosis (by age) of the innovations to (a) the transitory component of earnings (top) and (b) the persistent component of earnings (bottom). The dotted bands represent bootstrapped 95% confidence intervals.	61
15	Standard deviations of persistent shocks by previous earnings (left) and nonlinear persistence of the persistent component, by quantile of previous earnings and quantile of shock received in the current period (right). The dotted bands and transparent mesh represent bootstrapped 95% confidence intervals.	62
16	Cross-sectional variance of log earnings and log consumption by age. See Appendix 2.2 for details on their computation.	66
17	Growth in the cross-sectional variance of log consumption, data and implications of two earnings processes.	68
18	Growth in the cross-sectional variance of log consumption. Contribution of (1) age-dependence, (2) non-normality and (3) non-linearity to the growth of the cross-sectional variance of log consumption.	69

19	Cross-sectional average wealth (left) and variance of wealth holdings (right), by age and earnings process.	71
20	Persistence of the persistent component by quantile of previous earnings and current shock: NL (dark surface) vs non-normal age-dependent process (transparent surface).	72
21	Partial insurance coefficients on persistent shocks, ψ_t^p , by age . . .	76
22	Changes in the earnings distribution over the generations. Top: household heads; bottom: household earnings. Left: median earnings, center: average earnings, right: standard deviation of the log earnings distribution. PSID data.	89
23	Standard deviation of log earnings changes, by cohort	90
24	Higher order moments of log earnings changes, by cohort	90
25	Evolution of asset prices and returns. Left: Price to Income ratios for housing (PSID data). Right: Stock returns, S&P 500. Shaded areas correspond to NBER recessions.	91
26	Homeownership (left) and stock market participation (right), three generations. PSID and SCF data.	93
27	Nonlinear persistence, by cohort. Top left: 1940s; top right: 1960s; bottom: 1980s. For all cohorts: all agents below 40 years old. Figure C.8 in Appendix 3.3.5 provides the version with all ages included but thus varying composition across cohorts.	97
28	Average expected change in earnings (left), and Kelley's skewness of earnings changes (right), by previous earnings percentile and aggregate state of the economy. Percentiles refer to the distribution during an expansion.	98
29	Earnings by age with respect to the counterfactual in which a recession never occurs. Left: average earnings, recession at age 44. Middle: by initial earnings percentile, recession at age 44. Right: by initial earnings percentile, comparing ages 30 and 44.	99
30	Stock market returns, housing median price-to-income ratio	110
31	Life-cycle profiles for the 1940s cohort. Top left: homeownership by age; top right: proportion of households buying a house by age; bottom left: stock market participation; bottom right: percentage of all households with a mortgage by age.	118
32	Bottom: portfolio shares of assets by wealth decile at retirement age (left: PSID data, right: model).	118

33	Homeownership by cohorts, data vs. model	120
34	Homeownership by cohorts, by percentile of the earnings distribution at age 35	123
35	Homeownership by cohorts, the role of earnings risk vs earnings inequality. Left: 1960s generation; right: 1980s generation.	124
36	Stock market participation by age and cohort, data vs model. Top left: main model; top right: constant participation costs; right: constant participation costs, stocks with 401(k) tax properties	126
37	Homeownership by cohorts, benchmark vs. counterfactual earnings processes, empirical housing supply elasticity. Left: 1960s generation, right: 1980s generation.	127
38	Household responses to shocks. Top left: Blundell, Pistaferri and Preston (2008) coefficients. Top right: average net worth. Top bottom: average marginal propensity to consume, by age and cohort. Bottom right: including case in which household portfolios are fixed to their 1940s value.	129
39	Effects of shocks: left, consumption; right: homeownership.	131
40	Homeownership: projecting the 1980s cohort into the future	132
41	Wealth accumulation: projecting the 1980s cohort into the future. Top: constant stock market participation costs across cohorts; bottom: reduction in stock market participation costs for 1980s cohort. Units are multiples of average income.	133
42	Net worth by wealth percentile, ages 30-40, by generation. SCF data. Units are multiples of average income. For clarity, the top 5% and bottom 15% of the wealth distribution are not reported.	134
43	Homeownership by cohorts, data vs. model, canonical earnings process	135
A.1	Dutch male wages, computed using actual hours worked. Wage persistence (top left) and following moments of wage changes: standard deviation (top right), skewness (middle left), Kelley's skewness (middle right), kurtosis (bottom left), and Crow-Siddiqui kurtosis (middle right), by age group and previous earnings percentile.	155

A.2	The Dutch data: Non-robust measures of Skewness (left) and Kurtosis (right): male wages (first row) male earnings (second row) before-tax household income (third row) after-tax household income (fourth row).	156
A.3	The Dutch data: Male earnings by recent earnings.	157
A.4	The Dutch data: household earnings (left) and household pre-tax income (right)	158
B.1	Cross-sectional variance of log earnings over the life cycle, cohort effects vs year effects	164
B.2	Kelly skewness and Crow-Siddiqui kurtosis of male pre-tax earnings growth in the W2 (top panel) and PSID (central panel), and of household after-tax earnings growth in the PSID (bottom panel)	165
C.1	Median earnings (top) and average earnings (bottom), for household heads (left) and households (right), 2013 dollars, sample including SEO and immigrants.	174
C.2	Standard deviation of log earnings (left) and earnings changes (right), sample including SEO and immigrants	175
C.3	Homeownership, by cohort, sample including SEO and immigrants	175
C.4	Median (top) and average (bottom) earnings, 2013 dollars. Left: household heads, right: households. PCE-deflated sample.	177
C.5	Changes in the earnings distribution over the generations. Top: household earnings; bottom: household earnings for married households. Left: median earnings; center, average earnings; right: standard deviation of the log earnings distribution.	178
C.6	Homeownership by cohorts, PSID data. Left, sample restricted to married couples; right, sample restricted to households with at least one child.	178
C.7	Variance of log earnings over the life cycle, PSID data vs model-implied	179
C.8	Nonlinear persistence, by cohort. Top left: 1940s; top right: 1960s; bottom: 1980s. All ages	179
C.9	Direct stock market participation, three generations. SCF data.	180
C.10	Average OECD equivalence scale, by cohort	180
C.11	Homeownership by cohorts, benchmark vs. counterfactual earnings processes. Assuming fully inelastic housing supply.	183

C.12 Variance of earnings over the life cycle. Sold lines: data; dashed lines: NL process; dash-dot lines: canonical process	184
C.13 Stock market participation for the 1940s cohort	185
C.14 Homeownership by cohorts, data vs. model, relaxation of financial constraints in the early 2000s. Model with zero initial wealth.	186
C.15 Stock market participation by age and cohort, data vs model. Left: constant participation costs across generations; right: reduction of participation costs across generations. Model with zero initial wealth.	186
C.16 Empirical initial wealth. Left: density of wealth of 20-25 year olds, SCF data, measured as multiples of average income. Zero includes zero or negative. Right: model-implied homeownership with empirical initial wealth	188
C.17 Wealth accumulation: projecting the 1980s cohort into the future. With constant participation costs, comparing different reductions in average stock returns	188
C.18 Homeownership by cohorts, married people only.	189
C.19 Stock market participation over the life-cycle, by cohorts, case with calibrated per-period participation costs	190
C.20 Homeownership by cohorts, local correlation of house price shocks with income. Solid line with crosses: realizations + expectation, dotted line: expectations only	192
C.21 Alternative house sizes, main experiments	193
C.22 Percentage of households owning the small house by age. Left, robustness check in which it's smaller; right: main results	193

Introduction

Wage and earnings risk affect important economic decisions, including consumption, labor supply, savings, homeownership, and portfolio allocations. The extent to which people can self-insure against these risks is key to understand how people consume and save, but also to measure the scope for government policies that alleviate those risks.

Additionally, key features of earnings risk vary over the position of an individual in the income distribution and have changed over time. Understanding these changes is crucial to explain many structural transformations in the economy.

In Chapter 1, I quantitatively measure very rich features of earnings risk in Dutch administrative data. Namely, I study the characteristics of the distributions of the changes to male wages, male earnings, household earnings, pre-tax household income, and after-tax household income. I find that earnings display rich dynamics. These include age-varying persistence and variance of earnings changes, differences in earnings risk across the earnings distribution, and non-normal shocks. I decompose the sources of these features between hours and wages, and find that changes in hours are quantitatively more relevant to explain these nonlinear features. I then turn to studying the extent to which family labor supply and government insurance reduce these risks. I find that family insurance helps to smooth out large shocks, but government taxes and transfers play a much more important role in reducing these non-linearities.

In Chapter 2, I show that these rich features of earnings dynamics are also present in US administrative and survey data. I then estimate two processes for household after-tax earnings and study their implications within a standard life-cycle model. For both processes, I allow for a persistent and transitory component but while the first one is the canonical linear process with stationary shocks, the second one has substantially richer earnings dynamics, allowing for age-dependence of moments, non-normality, and nonlinearity in previous earnings and age. Allowing for richer earnings dynamics implies a substantially better fit of the evolution of cross-sectional consumption inequality over the life cycle and of the individual-level degree of consumption insurance against persistent earnings shocks. The richer earnings process also implies lower welfare costs of earnings risk.

Finally, in Chapter 3, I study how these features of earnings dynamics have

changed for different generations and their impact on homeownership and household portfolio composition. I find that earnings are riskier and more unequal for households born in the 1960s and 1980s than for those born in the 1940s. At the same time, despite the improvements in financial conditions that made it easier to borrow, younger generations are less likely to be living in their own homes than older generations at the same age. By using a rich life-cycle model with housing and portfolio choice that includes flexible earnings risk and aggregate asset price risk, I show that changes in earnings dynamics account for a large part of the reduction in homeownership across these generations. Lower-income households find it harder to buy housing, and some households delay homebuying decisions because their income is more unstable. As a result, they also accumulate less wealth. Relatively looser borrowing constraints help to explain how the 1980s cohort bought houses in a context of risky earnings and high house prices, and the reduction in the cost of access to financial markets can explain the intergenerational increase in stock market participation.

Chapter 1 Nonlinear Dynamics of Earnings in the Netherlands: Sources and Insurance Mechanisms

1.1 Introduction

Individual earnings are volatile. They are subject to infrequent, but relatively sizeable shocks of an asymmetric nature: large negative earnings shocks are more likely than large positive earnings shocks. Additionally, the persistence and variance of these shocks varies significantly over the life-cycle and over the distribution of income.

In this Chapter we study the size and distribution of earnings shocks in the Netherlands, document the source of these earnings fluctuations, and evaluate the role of insurance mechanisms against them. We start by documenting the distribution of wage shocks at the individual level by analyzing distributional measures of wage changes, including the standard deviation, skewness, kurtosis, and persistence, by age and previous earnings. To understand the role of individual-level labor supply and temporary unemployment spells, we compare the distribution of individual wage shocks with that of individual-level earnings. To analyze the role of family insurance through the labor supply of both partners, we compare the distributions of individual-level and household-level earnings. To examine the role of government insurance, we compare the distribution of household income, pre- and post-taxes, and transfers, by age group and previous earnings.

Our high-quality administrative data on income, taxes, and government transfers on individuals and households for the Netherlands (IPO) enables us to get precise estimates of the dynamics of wage shocks and the role of private and public insurance mechanism to mitigate these shocks. The results in this paper show clear evidence of non-linearity and age dependence of earnings dynamics, with high earnings risk for the people with lowest and highest earnings. Furthermore, wage and earnings persistence increase with age.

Our contribution to the literature is twofold. First, whereas previous studies mainly investigated shocks in individual earnings, we distinguish between shocks in wages and changes in hours worked (in line with [Hoffman and Malacrino \(2019\)](#)). As both may have different dynamics, this provides us with a better

understanding of the nature of income risk. Using high-quality administrative data on hours worked (DPA), we find that most of the fluctuations in earnings are related to changes in hours rather than changes in wages.

Second, we investigate the degree of insurance provided by spousal labor supply (by comparing individual earnings and total earnings at the household level) and insurance provided by the tax and transfer system (by comparing pre- and after-taxes household income). We find that the family is a relevant source of insurance in the Netherlands, but most of this insurance comes from income pooling rather than labor supply reactions of secondary earners or added worker effects. Taxes and, particularly, the transfer system play a much larger role in the reduction of income risk.

This Chapter contributes to a growing literature on higher-order moments of income shocks. [Güvenen, Karahan, Ozkan and Song \(2016\)](#) investigate higher order earnings risk using U.S. Social Security administrative data. They find substantial nonlinearities and non-normalities, but they can only study gross individual earnings process, so they cannot separate hours and wages or study additional insurance mechanisms.

Closely related to our study are a set of recent papers that use survey and administrative data from several countries to understand the drivers of earnings risk and the extent to which the household and the government insure against them. [Hoffman and Malacrino \(2019\)](#) use Italian administrative data to decompose earnings growth in changes in employment time and changes in weekly earnings. Like us, they find that changes in employment time are an important driver of earnings growth. [Halvorsen, Holter, Ozkan and Storesletten \(2019\)](#), using Norwegian data, and [Busch, Domeij, Güvenen and Madera \(2018\)](#), using data for the U.S. and Germany, attribute changes in earnings particularly to changes in wages. These international differences suggest that the institutional frameworks are important to determine whether wages or hours are the key margins of adjustment.

Similarly to our results, [Busch et al. \(2018\)](#) and [Halvorsen et al. \(2019\)](#) find that the benefit system is particularly important to insure workers against earnings fluctuations. Family insurance also matters, but to a lower extent. [Pruitt and Turner \(2018\)](#), on the other hand, use administrative data from the US. and find that the probability of the spouse entering employment rises when the male experiences earnings losses.

The remainder of this Chapter proceeds as follows. Section [1.2](#) describes

our data and approach, after which Section 1.3 presents the results. Section 1.4 concludes.

1.2 Data and approach

This section describes the data, our sample selection criteria, our wage and income measures, and the statistics that we set out to analyze.

Data sources We use administrative tax records from the Dutch Income Panel Study (IPO) and administrative data from Dutch payroll administrations (DPA).

The IPO data set contains a representative 1% population sample (of about 95,000 individuals) and their household members. The sample is randomly selected by Statistics Netherlands based on their national security number and is followed over time since 1989. Detailed information is available on, amongst others, personal income, household income, demographics, and labor market status. Because of a major tax reform, some of the income definitions in IPO changed in 2001. Our sample therefore starts in 2001.

The IPO data set has several important advantages over the use of survey data. First, the data is often collected from or checked with a third party. For instance, income measures are derived from tax records complemented with information provided by banks and other financial institutions. In addition, Statistics Netherlands performs several checks on the data to guarantee its quality. This drastically reduces or even eliminates measurement error and errors due to non-reporting. Second, individuals are followed for as long as they are residing in the Netherlands (as of December 31 of the sample year). We thus have little to no endogenous panel attrition. Panel attrition only occurs as a result of migration or death. New panel members enter the panel for the first time in the year of their birth, and immigrants to the Netherlands in the year of their arrival. Third, and very importantly, the IPO data set contains a detailed decomposition of labor and asset income, taxes and social insurance premia paid, and government transfers received for all household members. It also contains a detailed transfers breakdown, including unemployment insurance, disability insurance and social assistance. These features of the data allow us to measure the value of both family and government insurance.

The DPA payroll data is available since 2001 and is reported directly by employers to the tax authorities. It provides yearly information on the number

of days that a worker was employed and the number of hours worked on a typical working week, reported as a fraction of a full time job according to the sectorial workers' collective labor agreement, the so-called part-time factor. The part-time factor is not only based on contractual hours, but also on paid overtime hours.¹ Paid leave of absence, such as sick leave or parental leave, is counted in the data as hours effectively worked as long as wages are not reduced, and thus we cannot separate it out.²

Sample selection We select a sample of male earners age 25 to 60 to abstract from education and retirement decisions. We exclude self-employed workers³ and individuals with a very low attachment to the labor market. We include individuals with labor earnings of at least 2720 dollar a year (2200 euro) in 2014 prices. We equalize all measures of earnings that pertain to the entire household using the equivalence scale provided by Statistics Netherlands.

Variable definitions We study individual gross earnings, household gross earnings, household pre-tax (primary) income, household after-tax (disposable) income, and individual gross wages.

We define individual gross earnings as the total amount received by a worker in a given year according to their contracts, which includes employee's contributions to social security.⁴

We compute household gross earnings by aggregating individual earnings of all household members. By adding income from savings we obtain household pre-tax income. Finally, household after-tax income equals household pre-tax income minus income taxes plus allowances (healthcare, rent, child and child-care, study costs, and alimony) and transfers. Transfers are the sum of unemployment benefits, disability benefits, social assistance and pension benefits.

We compute individual gross wages w_{it} by dividing individual gross earnings y_{it} by a measure of hours worked h_{it} (Equation 1). Our measure of gross wages

¹For those workers whose contracts do not specify the explicit number of hours (such as zero-hours, min-max, or piece-rate pay contracts), there is information on the actual amount of hours paid.

²The part-time factor does not include overtime hours for full-time workers or overtime hours that are paid at a higher rate than usual hours. Fortunately, starting from 2006 we have detailed monthly information on hours that does include a very good measure of overtime. In Appendix 1.1 we show that considering that restricting the analysis to 2006 onwards and only considering this alternative, richer measure does not affect our main conclusions.

³That is, those for whom income out of self-employment is their main income source following [Güvenen et al. \(2016\)](#).

⁴In the Netherlands, these include a contribution for health insurance and a premium for unemployment, disability, and pension benefits.

w_{it} represents the amount an individual would earn when working a fulltime year. We obtain our measure of hours worked by linking administrative payroll data from the DPA to our main IPO sample, thus properly accounting for time spent unemployed, part-time work, and overtime. More precisely, to obtain h_{it} we use two elements. The first is the fraction of the year that a worker was employed according to payroll data. The second is the fraction of a typical working week that a worker spent at work during the time when she was on a contract (part-time factor). By multiplying both we obtain the fraction of an effective working year that an individual spent at work (Equation 2).⁵

$$w_{it} = \frac{y_{it}}{h_{it}} \quad (1)$$

$$h_{it} = \frac{\text{Weeks worked per year}}{52} \times \frac{\text{Hours worked per week}}{\text{Usual working week}}. \quad (2)$$

Approach As standard in the literature, we purge age and time effects from log wages by running the following regression and identifying its estimated residuals as wage shocks

$$\log w_{it} = \beta_1 \text{age}_{it} + \beta_2 \text{age}_{it}^2 + \alpha_t + u_{it}. \quad (3)$$

The subscript i refers to an individual, t is year, α_t represents year fixed effects, and the error term u_{it} captures the stochastic component of wages.

Because the widespread modeling of wage shocks as an AR(1) process implies strong restrictions on wage changes that previous work has found to be violated in the data, rather than making this functional form assumption, we compute key moments of wage shock changes ($\hat{u}_{it} - \hat{u}_{it-1}$), including their standard deviation, skewness, and kurtosis.

These moments derive from interesting and important economic mechanisms. For instance, negative skewness can come from a job ladder model (see, for instance, [Graber and Lise \(2015\)](#)) in which people staying on the job experience small wage raises most of the time, but people losing their job often experience a large wage and earnings drop. This kind of model can also explain some kurtosis: most wage changes are small, but then there is a small fraction of people experiencing large wage changes, due to job loss, or job and career switches, for instance. In addition, the persistence of these wage changes might depend on one's age and current earnings. A young worker is more likely to

⁵This measure can then be transformed into actual hours by multiplying it by the number of working hours in a typical working year.

switch jobs and careers to figure out what he or she is best at, which tends to lower the persistence of their wage changes. An old worker might switch to a part-time or less demanding job, thus also having lower wage shocks persistence. Finally, earnings persistence might depend also on previous earnings. For instance, high earners might be experiencing more wage uncertainty than those with a middle-class income.

To measure skewness, we compute the conventional measure of skewness (Pearson’s or standardized third moment). Because this measure is very sensitive to outliers (deviations from the mean are cubed), we also compute the robust Kelley’s coefficient of skewness, which is given by

$$S_K = \frac{P_{90} + P_{10} - 2P_{50}}{P_{90} - P_{10}}, \quad (4)$$

where a zero implies a symmetric distribution, positive values represent right skewness, and negative values represents left skewness.

To measure kurtosis, we start with the conventional measure given by the fourth standardized moment, but we also compute the robust Crow-Siddiqui kurtosis which is given by

$$S_{CS} = \frac{P_{97.5} - P_{2.5}}{P_{75} - P_{25}}. \quad (5)$$

The term S_{CS} is large if $P_{97.5} - P_{2.5}$ is large relative to the probability mass that is concentrated between P_{75} and P_{25} , corresponding to heavy tails.

Finally, we analyze persistence by age, by regressing \hat{u}_{it+1} on \hat{u}_{it} for different ages.

To investigate insurance mechanisms, after studying wages, we repeat the analysis for individual-level earnings, household earnings, household pre-tax income (earnings and income from savings) and household after-tax (disposable) income. The comparison of wages and earnings is informative about self-insurance through labor supply. The comparison of individual-level and household-level earnings is informative about family insurance through the labor supply of the spouse. The comparison between household pre-tax income and household disposable income helps shed light on the role of insurance by the government through transfers and progressive taxation.

1.3 Results

In this section, we first discuss the properties of male wage changes. Then, we compare them with those of male earnings, household earnings, pre-tax household income and disposable household income. We also discuss what they imply in terms of family and government insurance.

1.3.1 Male wages

The top left panel of Figure 1 displays wage persistence by age and shows large age variation in male wage persistence, unlike typically assumed by a standard AR(1) process. Wage persistence starts from a low value of 0.65 at age 25, consistently with younger people switching jobs and careers to figure out what job is the best fit for them. Many of them also have temporary contracts. It then increases fast, reaching 0.85 at age 30, and 0.9 at the age of 40. After that age, it remains flat. This is good news for older workers with high wages, but low wage workers may become more vulnerable when this becomes a persistent situation. At younger ages, on the other hand, bad (and good) wage shocks are not as long-lived as they would be if their persistence were much higher, as assumed by many models evaluating policy interventions and taxation.

The top right hand side of Figure 1 displays the standard deviation of wage changes by age group and previous earnings.⁶ Several features of the data are worth noticing. First, the variability of wage changes is about 1.5 times larger at the lowest percentiles of previous earnings (0.35) than for workers in the middle percentiles of previous earnings (0.23). Second, this variability increases somewhat (from 0.23 to almost 0.25) for previous earnings above the 90th percentile. Third, workers in the youngest (25-34) and the oldest (55-59) age group with previous earnings below the 30th percentile experience the largest wage change volatility.

The differential patterns for young workers may be related to flexible contracts. In contrast, workers at late stages of their career might have a higher prevalence of sick-leave due to longer-lived health problems. In the first two year of sick-leave wages may decline, as the employer is required to continue paying at least 70% of wages in the first two years of sick leave. Thereafter, one may become eligible for disability benefits.

⁶In all figures we use the same horizontal axis, which is the distribution of previous earnings. Using previous wages on the horizontal axis of Figure 1 does not change the results.

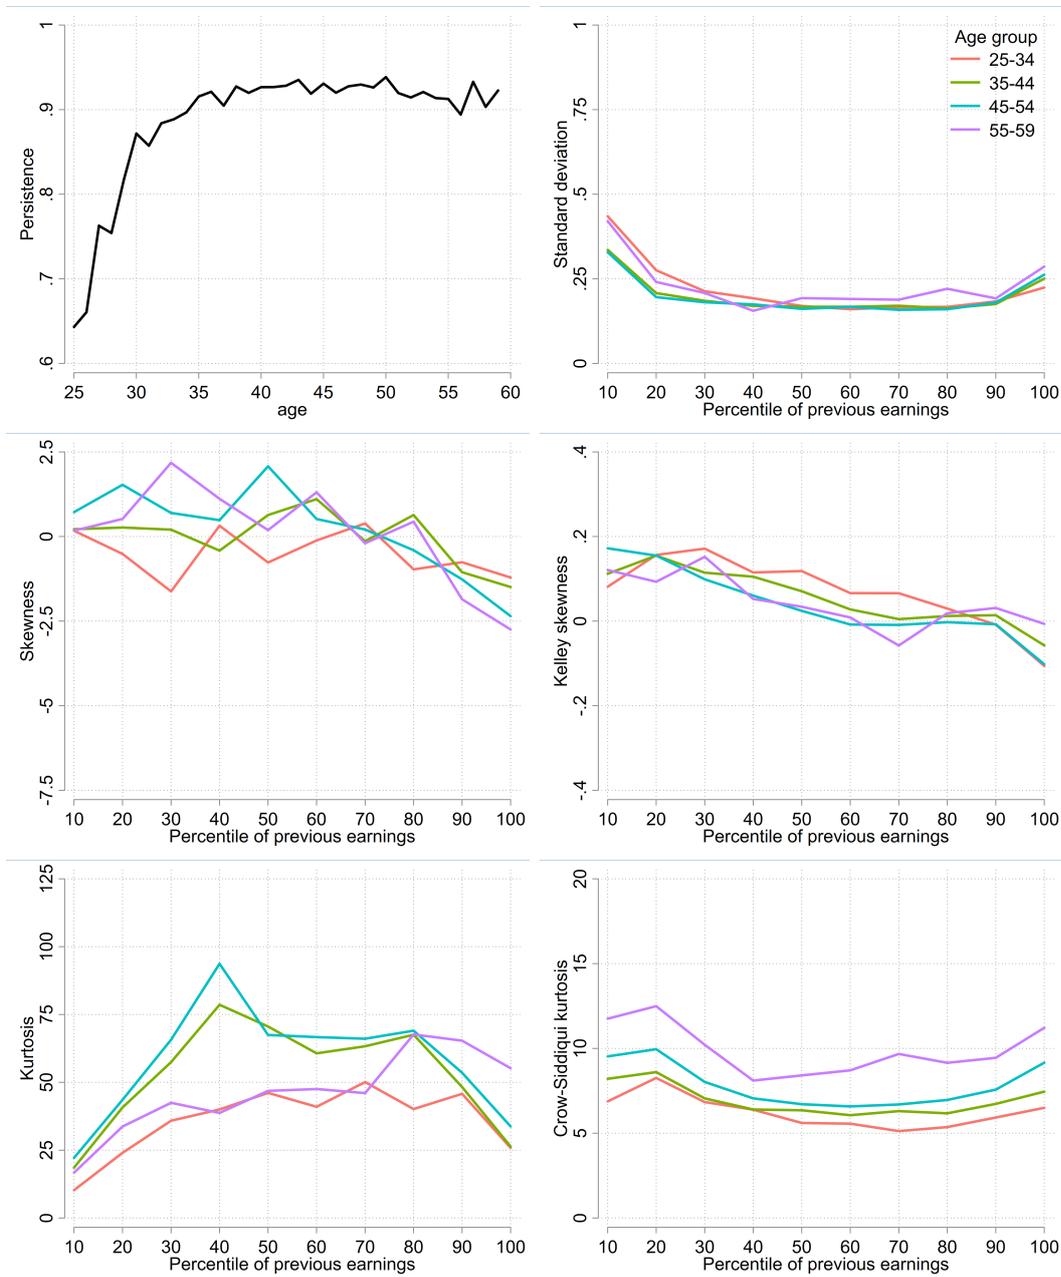


Figure 1: Dutch male wages. Wage persistence (top left) and moments of wage changes: standard deviation (top right), skewness (middle left), Kelley’s skewness (middle right), kurtosis (bottom left), and Crow-Siddiqui kurtosis (middle right), by age group and previous earnings percentile.

High volatility, however, may be due to both upward or downward mobility. To study asymmetries in wage changes, the middle left panel of Figure 1 plots a measure of skewness. Wage skewness starts around zero at low levels of previous earnings. After the 70th percentile of previous earnings skewness becomes negative, reaching -2.5 for the older age groups with the largest previous earnings. This means that workers with higher previous earnings are more likely to experience a relatively large negative wage shock, rather than a wage increase. Also, young workers with previous earnings below the 50th percentile, have more negative skewness than their older counterparts. They may experience relative large negative wage changes after an unemployment spell.

Because this conventional measure of skewness is quite sensitive to the tails of the distribution, we also report (middle right panel of Figure 1) Kelley's skewness, which is robust to outliers. Once we eliminate unusually negative wage shocks, skewness is zero or slightly positive for most age groups and most of the wage distribution. Between the 20th and 80th percentile of previous earnings, the youngest age group has the most positive skewness, which can be related to promotions early in people's careers. Looking at both measures jointly, we find little evidence of negative skewness in Dutch wages, with the only exception of the very income-richest.

The bottom left panel of Figure 1 shows the kurtosis of wage changes. It is increasing at the bottom of previous earnings (up to the 40th percentile). Workers in their prime working lives (aged 35-54) have the highest kurtosis and thus face the distribution of wage changes with the fattest tails. Broadly, this suggests that wage shocks are very infrequent but that, when they happen, they tend to be of a large magnitude.

The Crow-Siddiqui kurtosis measure (bottom right panel), which is robust to outliers, confirms this intuition. Although it displays less variation over the distribution of previous earnings, its age patterns are starker. Younger workers experience relatively smaller but more frequent wage shocks, whereas for older workers these shocks become larger but much less prevalent. This may be due to a lower prevalence of flexible contracts and higher employment protection for older age groups, which are likely to make wages of older and higher income workers more rigid.⁷

Thus, Figure 1 shows strong evidence in favor of age-variation, non-linearity

⁷By law it is very difficult to lower someone's wage rate. At the end of working life wage scale ceilings can be restrictive, especially in the public sector [Deelen and Euwals \(2014\)](#).

and non-normality of wage changes: wage persistence is the lowest for the youngest, wages are more variable for the highest and lowest earnings, and there are hardly any shocks for most and large shocks for some. However, skewness is very close to zero for most of the earnings distribution, suggesting an almost symmetrical distribution of wage changes for most, with the exception of the highest earners.

1.3.2 Male earnings

Previous studies have focused mostly on the features of earnings shocks. Our data, with a very precise measure of hours, is particularly well suited to understand whether these come from shocks in wages or hours. In Figure 2 we compare the statistics for male wages (left panel) with those for male earnings (right panel). There are several economically relevant differences.

First, looking at persistence (top panel), we observe that it increases more rapidly by age for wages than for earnings. With respect to standard deviation (second panel), we find that at the bottom of the distribution of previous earnings, the variability in earnings changes is larger than the variability of wage changes. This suggests that there is important action at the hours margin, and that much of the earnings fluctuations of lower earners are due to temporary unemployment, reductions in working time, or labor supply decisions.

The third panel shows that Kelley's skewness of earnings changes is in general more negative than that of wage changes, particularly for the very oldest. The picture for centered skewness is even starker (fourth panel). Skewness of earnings is much larger, in absolute value, than skewness of wages. This suggests that the driver of most large negative changes in earnings is a reduction in hours (i.e., unemployment, temporary disability, etc.), with a much more limited role for wages, with the only exception of the very highest earnings.

For the oldest age group the variability in earnings changes is slightly higher than the variability in wage changes, the skewness is more negative, and the kurtosis (bottom panel) is more positive across the whole distribution of previous earnings. This pattern could be explained by a higher prevalence of job-loss, but also by early retirement and a (voluntary) reduction in working hours among the oldest age group.

To further understand the role of hours and wages in explaining the rich nonlinear patterns in male earnings, Figures 3 and 4 decompose their relative contributions to three main measures: the variance, skewness, and kurtosis of

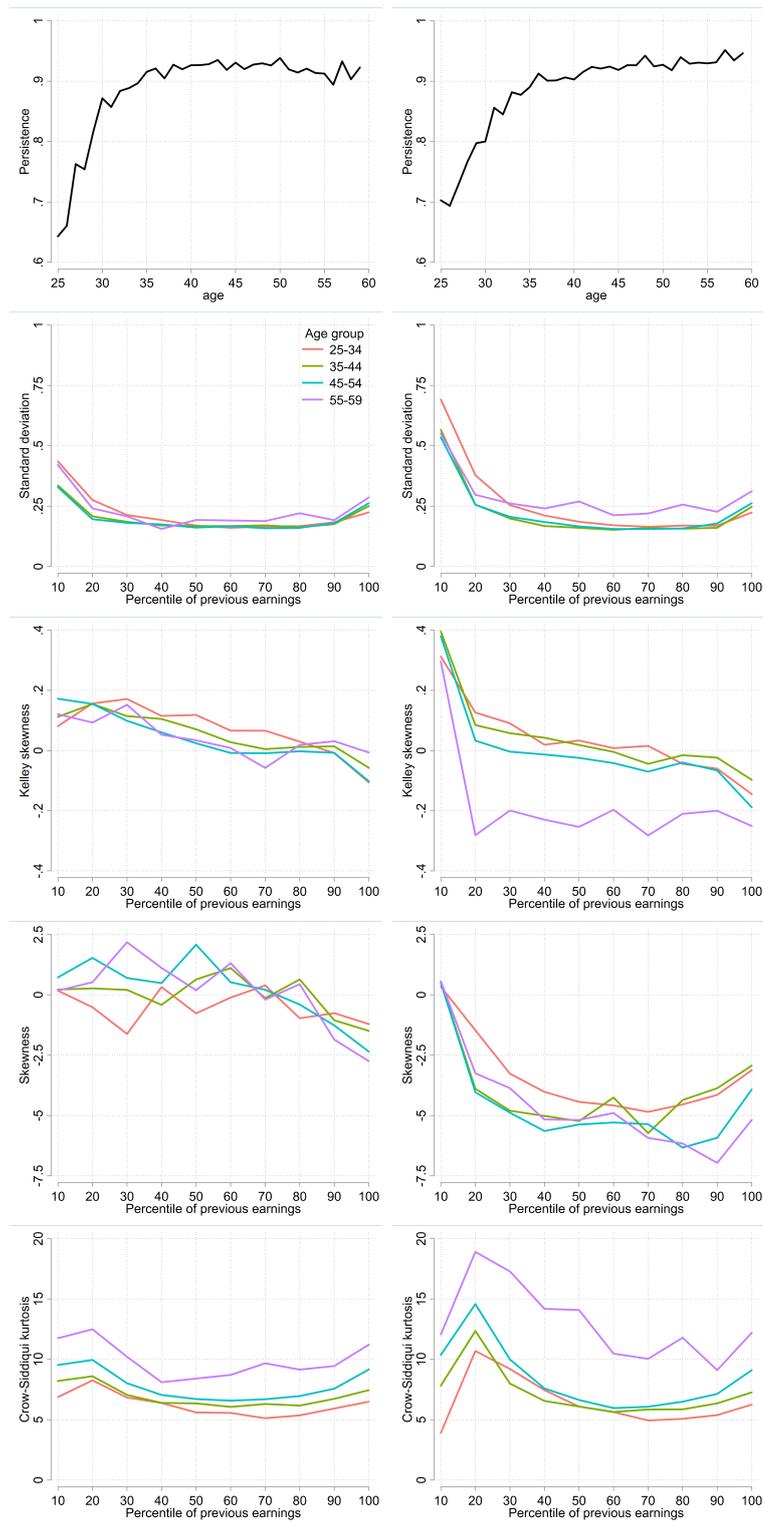


Figure 2: NL, male wages (left) and male earnings (right). Persistence (top row), standard deviation (second row), Kelley's skewness (third row), skewness (fourth row), and Crow-Siddiqui kurtosis (bottom row).

male earnings changes.

In Figure 3 we observe that, for most households, the variance of wages and hours worked is very low. It is larger for the lowest earners, for whom most of the fluctuations are related to hours rather than wages, and for the highest earners, for whom the opposite is true.

The two bottom panels study the contributions to the variance separately for those individuals who have suffered a positive change in hours (left) and those who have suffered a negative change in hours (right). For those whose earnings go up, with the exception of the bottom tail of the earnings distribution, the variance of hours is almost zero. Thus, for most of the income distribution increases in earnings come through increases in wages at a constant amount of hours worked. The covariance between hours and wages is very low, which suggests that there are almost no labor supply reactions to wage changes at these ranges of earnings in which most individuals are in full-time employment already.

The lowest earners present remarkably different patterns. For many of them, the cause of their previously low earnings was a low amount of hours worked in a year, and most of their positive shocks are driven by increases in hours worked rather than in wages. Indeed, we observe a negative correlation between hours and wages at this earnings levels, which may reflect a wage penalty when re-entering the labor market after experiencing an employment shock.

For those whose earnings go down, variations in earnings are mostly associated with variations in hours. Thus, negative shocks to earnings seem to be mostly driven by changes in hours.

This intuition is further confirmed in the left panel of Figure 4. In the Netherlands, most of the negative skewness in earnings is driven by changes in hours rather than wages. Thus, the relatively large earnings skewness mostly reflects temporary periods out of employment or reductions in the numbers of hours worked per week. This is consistent with the evidence presented in Hoffman and Malacrino (2019) for Italy, but at odds with the findings in Busch et al. (2018) for Germany. These international differences suggest that the institutional framework that governs the labor market is crucial to determine the sources of earnings fluctuations and whether adjustments occur at the margin of hours or wages.

Co-skewness measures whether large changes in hours and wages occur at the same time. Whereas in Norway co-skewness plays a substantial role in

explaining negative skewness of earnings growth (Halvorsen et al., 2019), this is not the case for the Netherlands, where it is very close to zero. Positive kurtosis in the right-hand-side panel of Figure 4 is also somewhat more driven by hours than wages. Most individuals do not change working hours between one year and the other and this leads to a lot of relatively small changes in earnings.

Figure 5 further shows the extent to which changes in male earnings are related to changes in hours and wages. The horizontal axis measures changes in earnings and the vertical axis measures changes in hours or wages, both expressed in terms of log points. Each dot in the line represents a decile of changes in male earnings, and each of the graphs represent three different positions in the previous earnings distribution (first decile, median, and ninth decile, from left to right).

For instance, the leftmost datapoint in the left panel of Figure 5 shows that earners with low previous earnings (1st decile of the previous earnings distribution) and a very bad earnings change (1st decile of the change in male earnings) have suffered on average an 80% decrease in their earnings. Of these, above 70 percentage points are accounted for by a reduction hours, with a bit below 10 percentage points due to a reduction in wages.

Naturally, those with the lowest previous earnings are also those who experience the largest earnings increase in relative terms (140%, the rightmost datapoint in the left panel of Figure 5). For those, again, most of the change is due to an increase in hours, with a small role for wages. The largest share of people, however, face very small changes to their earnings, hours, and wages, and that's why many dots are located very close to zero.

These patterns are different for people who had around median earnings last period (central panel). For them, negative shocks are due to both changes in wages and hours, and positive shocks are almost entirely due to increasing wages. In the top of the distribution of previous earnings (right panel), negative shocks in earnings are mainly due to drops in wages but also due to changes in working hours. This observation agrees with the evidence in Figure 2, where we observed that there was negative skewness of wage growth for this earnings group. Positive shocks, on the other hand, are only due to wage growth, as most people in this earnings range were working full time in the previous period.

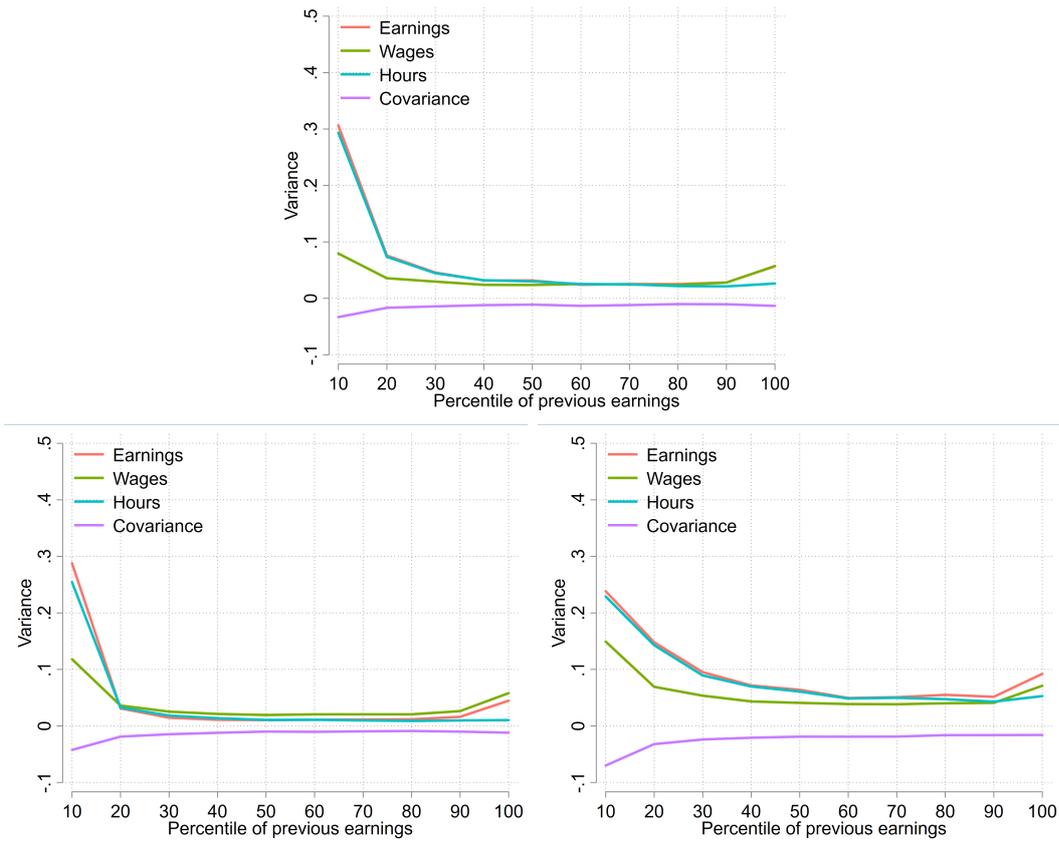


Figure 3: Variance of changes in male earnings, wages and hours. Top: all workers; bottom left: workers with a positive hours change; bottom right: workers with a negative hours change.

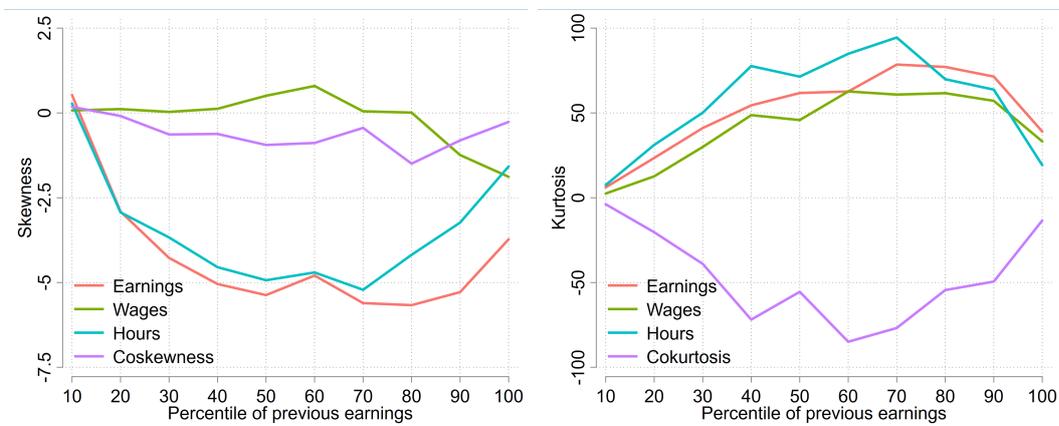


Figure 4: Skewness and kurtosis of changes in male earnings, wages and hours

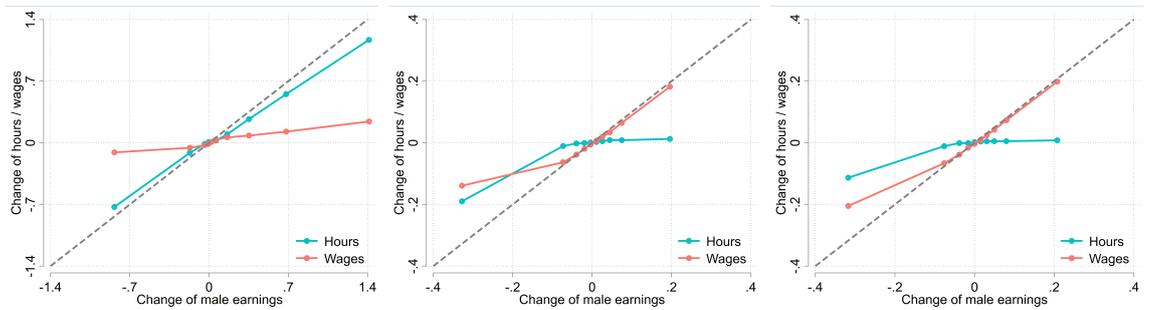


Figure 5: Male earnings changes versus hours and wage changes. Each dot represents a decile of changes in male earnings. First decile of previous earnings (left), median decile of previous earnings (middle), 9th decile of previous earnings (right).

1.3.3 Household insurance

Next, to investigate the effect of insurance within the household, we compare the nature of changes in male and household earnings (left vs. right panels of Figure 6). First, persistence is very similar for male and household earnings. Among older workers the standard deviation is a little bit lower and Kelley's skewness is somewhat less negative for changes in household earnings compared to male earnings. Interestingly, for younger workers we find higher standard deviations and more negative Kelley skewness for household earnings compared to male earnings, which could reflect female spouses who reduce working hours after the birth of children.

The bottom two panels of the figures show that the labor supply of the secondary earner plays an important role in reducing the skewness and kurtosis of household earnings earnings compared to that of male earnings and wages. This means that on the household level there are more frequent but small changes in earnings, compared to less frequent but larger changes in male earnings and wages. Thus, in the Netherlands the family has a role in reducing the risks that households face. However, these features of the data might be either due to a pooling of earnings within the household or due to an increase in the labor supply of women when their husbands experience a negative earnings shock (added worker effect).

Figure 7 examines which of these two different channels of within-household insurance is more prevalent in the Netherlands. It represents the average change of women's hours between years t and $t + 2$ as a response to changes in male earnings between t and $t + 1$ for couples. If there is an added worker effect, the number of hours worked by the woman in the household would respond to earnings shocks suffered by the man; by looking at two-year windows we can

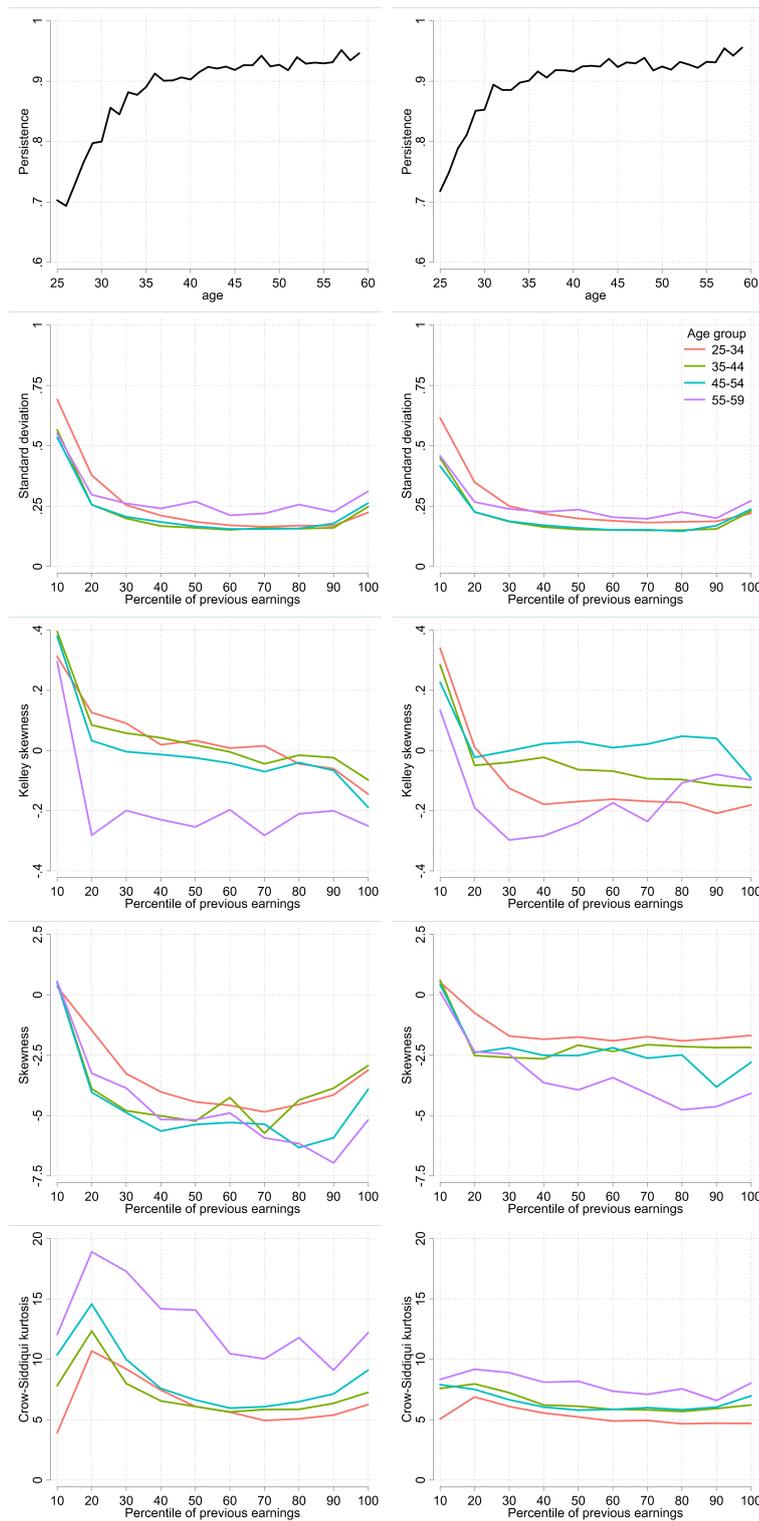


Figure 6: NL, household earnings (left) and disposable income (right). Persistence (top row), standard deviation (second row), Kelley's skewness (third row), skewness (fourth row), and Crow-Siddiqui kurtosis (bottom row).

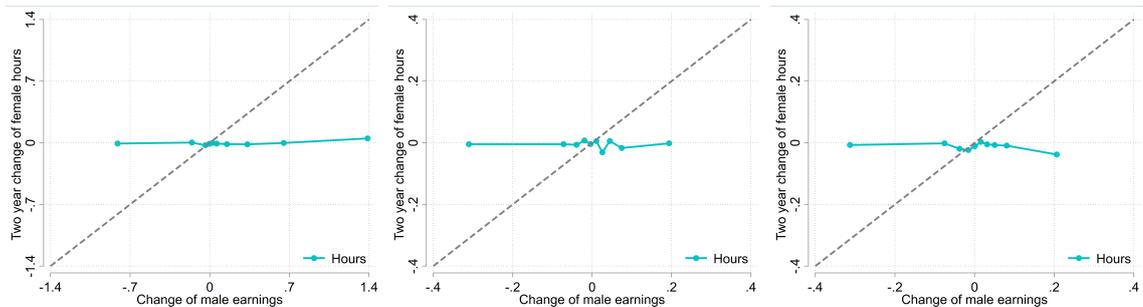


Figure 7: Male earnings changes and female labor supply. Each dot represents a decile of changes in male earnings. Lowest decile of previous male earnings (left), median decile of previous male earnings (center), 9th decile of previous male earnings (right)

capture changes in female labor supply which are not exactly contemporaneous to the man’s earnings shock. We do not find any association between changes in male earnings and changes in women’s hours worked, indicating that it is mostly income pooling which explains the reduced earnings risk that households face. This is in line with findings for Norway ([Halvorsen et al. \(2019\)](#)), and may be due to correlated labour market opportunities of spouses. The only case in which we find a noticeable labor supply reaction in the Netherlands is women who reduce hours worked as a response to large positive changes in male earnings for those who were already high earners.

1.3.4 Insurance from taxes and transfers

To investigate the role of government insurance, Figure 8 compares the properties of household pre-tax income (left panel) with disposable income (right panel) for the Netherlands.⁸

Taxes and transfers make a huge difference for the measures of risk that we focus on, especially at the lower end of the income distribution and for older households. For disposable income, the standard deviations are lower and both measures of skewness become less negative. For instance, the standard deviation of household income changes at the lowest percentiles of previous earnings declines from about 0.75 before taxes and transfers to a little over 0.37 after taxes and transfers. The reduction in the standard deviations and Kelley’s skewness is especially apparent for workers in the oldest age group. For them, Kelley’s skewness becomes almost zero. The Crow-Siddiqui kurtosis further drops from about 8 at the household level (it peaked at about 17 for wages and male earnings) to well below 7 after taxes and transfers.

⁸Household pre-tax income contains earnings and income from savings. In Appendix 1.4 we show that allowing for capital income makes little difference for household income dynamics.

Figure 9 summarizes the roles of household and government insurance by showing the pass-through of changes in male earnings to before- and after-tax income. It shows that taxes and transfers offset positive and negative changes in male earnings, especially for households at the bottom of the distribution of previous earnings. For example, households in the 20th percentile of previous earnings with a negative earnings shock of 60% experience on average a 40% drop in pre-tax household income, but only a 10% drop in disposable household income. Households in the 50th and 80th percentile of previous male earnings experience smaller changes in male earnings and receive, as expected, less insurance from progressive taxation and transfers (the difference between the slopes of the blue and the red lines is smaller).

Given that government insurance is especially prevalent in the Netherlands and especially so at older ages, Figure 10 further breaks down the role of various government programs for our 55-59 age group, by sequentially adding specific transfer programs or taxes. The graphs show that disability insurance greatly reduces the standard deviation of household earnings changes below the 20th percentile of previous earnings, while unemployment insurance generates a significant reduction even at higher levels of previous earnings. It also shows that, for this age group, early retirement transfers play a much larger role in reducing variation in household income than taxes. The bottom graph of Figure 10 shows that negative skewness is completely offset by taxes and transfers in the bottom 45 percentiles of the distribution of previous earnings, whereas it is partly offset between the 45th and the 80th percentile of previous earnings.

The analyses make clear that the government and private pensions provide a lot of insurance in the Netherlands. Progressive taxation reduces earnings variability and the benefit system (unemployment insurance, disability insurance, and welfare) and private pensions reduce income variability. In particular for older workers and for the bottom of the distribution of previous earnings, taxes and transfers effectively eliminate large negative shocks, such that negative skewness disappears and the kurtosis is reduced.

1.4 Conclusion

This Chapter studies the nature of income dynamics over the lifecycle. We use unique administrative data from the Netherlands to analyze shocks in wages and hours worked. Furthermore, we investigate the degree of insurance provided by spousal labor supply and by the tax and transfer system.

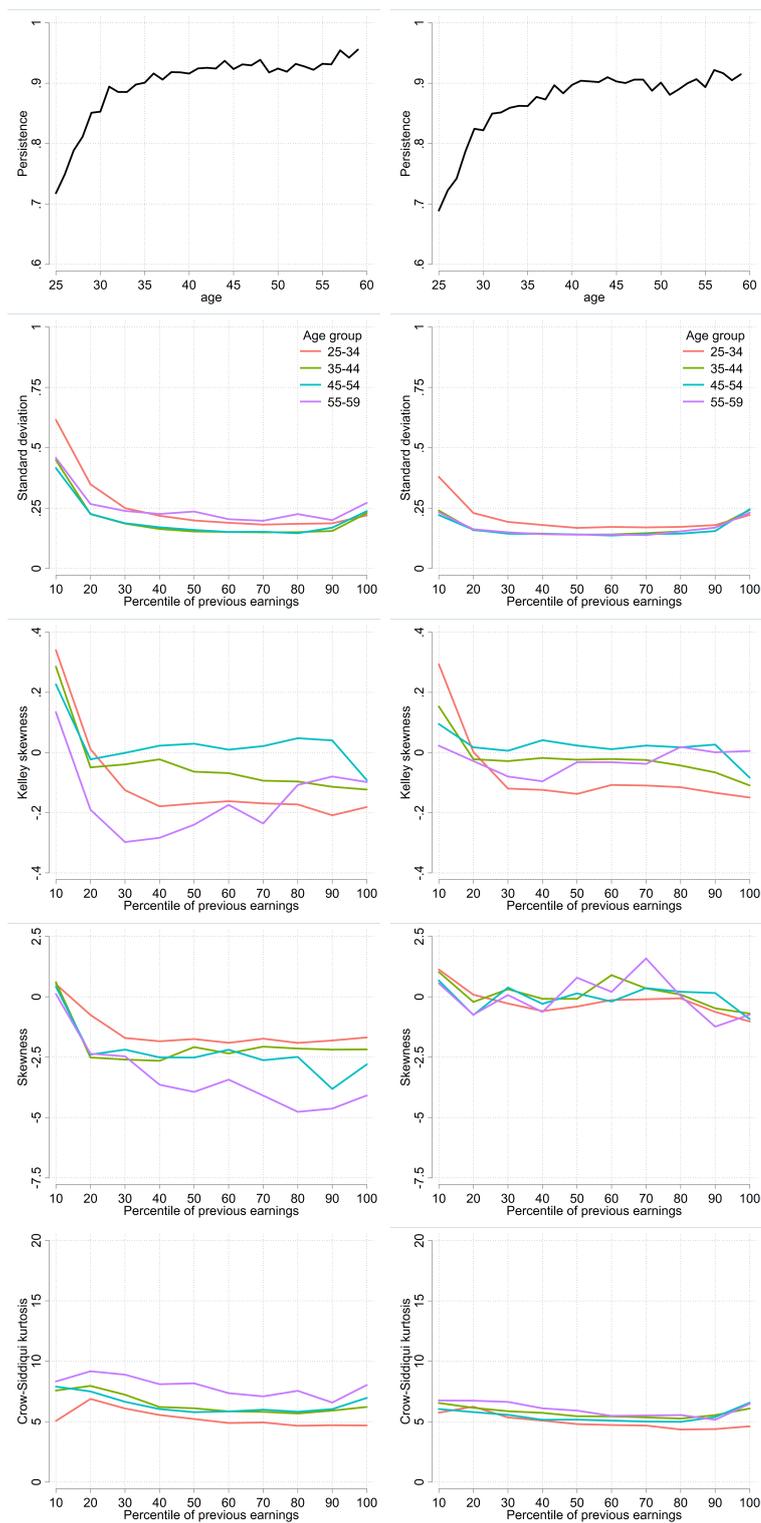


Figure 8: NL, household earnings (left) and disposable income (right). Persistence (top row), standard deviation (second row), Kelley's skewness (third row), skewness (fourth row), and Crow-Siddiqui kurtosis (bottom row).

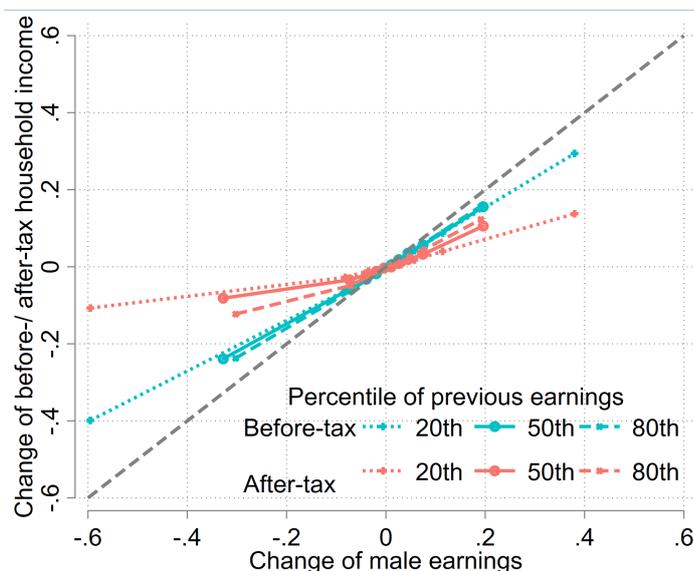


Figure 9: Household before- vs after-tax income. Each dot represents a decile of changes in male earnings.

The results show clear evidence of non-linearity and age dependence of wages and earnings in the Netherlands with higher earnings risk for the lowest and highest earners. Except for outliers, large wage shocks are mostly positive in the first half of the distribution and negative in the top decile of the distribution. In line with [Hoffman and Malacrino \(2019\)](#), we find that changes in employment time are an important driver of changes in earnings. More specifically, we show that this holds especially at the lower part of the earnings distribution. In the top, upward mobility is driven by positive wage shocks for workers who were mostly already working full time.

Our results show that large downward shocks in earnings are more likely than large upward shocks. Especially for older workers and above the lowest income group, people reach wage scale ceilings and exhaust opportunities to move up, while negative earnings risk due to sickness, long term unemployment and retirement increases. For most workers, however, earnings stay about the same from one year to the other (the kurtosis is high). Especially the wages and earnings of older workers appear to be rigid, apart from some outliers.

In the Netherlands women's earnings do not reduce the standard deviation of income risk at the household level. Indeed, for the age group 25-34 the variance even increases after the 30th percentile, probably due to the birth of children. However, income pooling within the household does substantially reduce skewness, thus suggesting that the presence of a secondary earner in the household can smooth out large negative shocks. We do not find evidence for

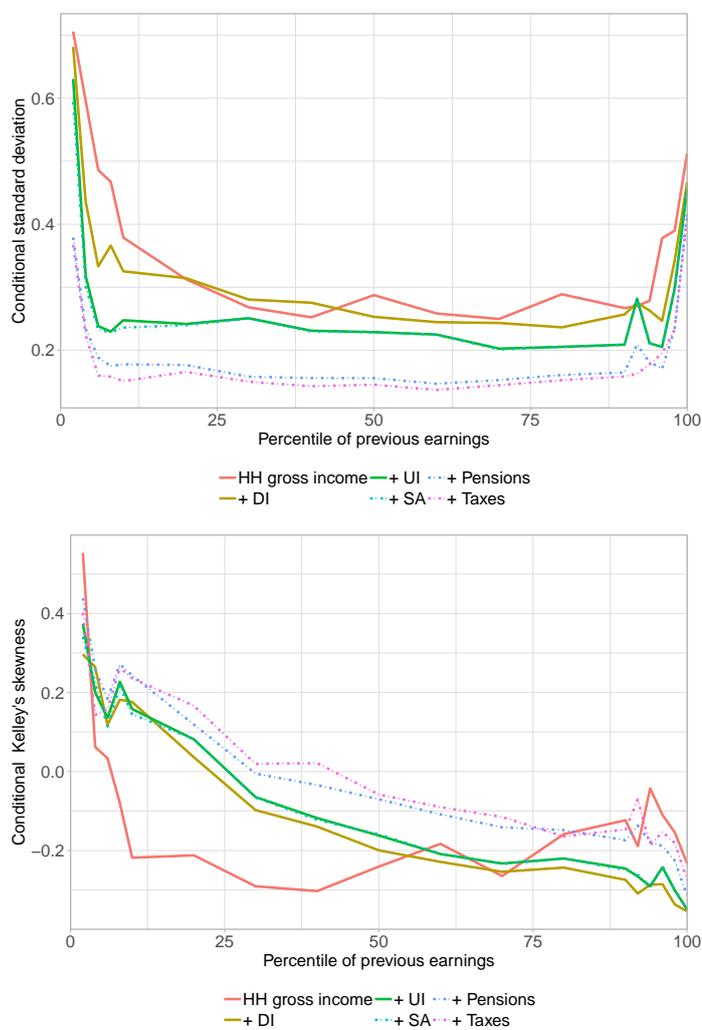


Figure 10: NL, age 55-59, Relative contribution of transfers and taxes to the standard deviation of household income. Red line, household gross income, gold line: including disability insurance, green line: also including unemployment insurance, dotted green line: also including social assistance, dotted blue line also including pensions, dotted red line: also net of taxes.

an added worker effect in the Netherlands.

Comparing family and government insurance we find that the government plays a much larger role in reducing earnings risk. A breakdown by government programs for older workers in the Netherlands shows that disability and unemployment insurance programs reduce income risk, especially for the lowest quarter of the male earnings distribution. Pensions and taxes (to a lower extent) reduce earnings risk across the whole distribution.

After gathering this empirical evidence regarding the dynamics of earnings in the Netherlands, we now turn to studying whether these patterns are also true for the United States and to evaluating to which extent richer earnings dynamics in disposable income affect household consumption decisions, self-insurance, and welfare.

Chapter 2 The Implications of Nonlinear Earnings Dynamics for Household Self-Insurance and Welfare

2.1 Introduction

Macroeconomic models with heterogeneous agents are ideal laboratory economies to quantitatively study a large set of issues that include household behavior under uncertainty, inequality, and the effects of taxes, transfers, and social insurance reforms.⁹ Earnings risk is a crucial source of heterogeneity in these models and its stochastic properties determine how saving and consumption adjust to buffer the impact of earnings shocks on current and future consumption. Appropriately capturing earnings risk is therefore important to understand consumption and wealth inequality, the welfare implications of income fluctuations, and the potential role for social insurance.

With few notable exceptions, most quantitative macroeconomic models adopt earnings processes that imply that persistence and other second and higher conditional moments are independent of age and earnings histories, and that shocks are normally distributed. The *canonical* permanent/transitory process is a popular example.

A growing body of empirical work, though, provides evidence that households' earnings dynamics feature non-normality, age-dependence, and nonlinearities, and devises flexible statistical models that allow for these features. Chapter 1 of this thesis contributed to this strand of the literature by looking at Dutch administrative data. Other recent work takes advantage of large, administrative datasets (e.g., W2 confidential Social Security Administration earnings data in [Guvenen et al., 2016](#)) and new methodologies applied to survey data sets like the Panel Study of Income Dynamics (PSID) ([Arellano, Blundell and Bonhomme, 2017](#)) to show that changes to pre-tax, individual male earnings display substantial skewness and kurtosis and that the persistence of shocks depends both on age and current earnings.¹⁰

⁹For instance, [Scholz, Seshadri and Khitatrakun \(2006\)](#) study the adequacy of savings at retirement, [Storesletten, Telmer and Yaron \(2004b\)](#); [Krueger and Perri \(2006\)](#); [Heathcote, Storesletten and Violante \(2010\)](#) study the evolution of consumption and [Castañeda, Díaz-Giménez and Ríos-Rull \(2003\)](#); [De Nardi \(2004\)](#); [Cagetti and De Nardi \(2009\)](#) study the evolution of wealth inequality over the life cycle, while [Conesa, Kitao and Krueger \(2009\)](#) study the optimal taxation of capital income.

¹⁰These features are consistent with several factors that affect the working lives of individuals. For

In line with the evidence provided in Chapter 1, in this Chapter we show that, in the United States, all of these rich dynamics are present in individual pre-tax earnings, both in the W2 tax data and the PSID, and in household, after-tax earnings, which are the relevant source of labor income risk at the household level.¹¹

Our main contribution is to analyze the effects of richer earnings dynamics on consumption, wealth, and welfare, both in the cross-section and over the life cycle. We use the econometric framework recently proposed by [Arellano et al. \(2017\)](#) to separately identify the distributions of the persistent and transitory components of earnings while allowing for non-normality of shocks, non-linear persistence, and, in general, a rich dependence of the two distributions on age and (in the case of the persistent component) previous earnings. We use PSID data on after-tax, household earnings to estimate two different earnings processes: a richer earnings process along the lines of [Arellano et al. \(2017\)](#) and a “canonical” linear earnings process with a persistent and transitory component and normal innovations, like the one used in [Storesletten et al. \(2004b\)](#). For each process, we use two Markov chains to approximate the conditional distributions, respectively, of persistent and transitory shocks and introduce them into a partial-equilibrium, life-cycle model with incomplete markets to compare their implications for consumption, wealth, and welfare.¹²

Our main findings are as follows. First, compared to the canonical earnings process, the richer earnings process better fits the observed evolution of consumption inequality over the life cycle. More specifically, under the canonical earnings process, the growth in the variance of consumption substantially overshoots its data counterpart at all ages, while our richer process generates a realistic profile up to ages 50-55, when early and partial retirement start being important. The improved fit is due to the rich features of the earnings data that we model and to the households’ precautionary saving response to them.

instance, younger people tend to change jobs more frequently and this implies that the persistence of their earnings is lower. In addition, for most workers, earnings vary little from year to year and shocks are infrequent but can be of large magnitude, such as job loss or a career change, when they happen. This is captured by the high level of kurtosis displayed by earnings changes.

¹¹One caveat is that, in line with much of the previous literature, we take earnings as a primitive while earnings fluctuations likely partly reflect endogenous labour supply choices.

¹²Although this study is highly indebted to [Arellano et al. \(2017\)](#), it differs from it in important respects. First, [Arellano et al. \(2017\)](#) estimate their earnings process on *pre-tax* earnings of households headed by *participating and married males* while we use *after-tax* earnings for *all* households. Secondly, [Arellano et al. \(2017\)](#) estimate the consumption rule semi-parametrically while we obtain it using the full model structure and, for the same reason, we can study welfare implications. Finally, the emphasis of [Arellano et al. \(2017\)](#) is on the consumption response to earnings shocks while we consider a wider range of implications.

In particular, the age-dependent persistence and variance of earnings innovations account for the main share of the improvement of the fit between age 25 and 45, while non-normality and, in particular, nonlinearity (for instance, the fact that persistence varies with the level of previous earnings) drive the improvement between age 45 and 55.

An alternative, and possibly more intuitive, measure of self-insurance is the extent of consumption passthrough of shocks to disposable earnings onto consumption. Our second finding is that the richer earnings process implies a consumption passthrough of persistent earnings shocks broadly consistent with the data. Its value is 0.57 which is within one standard deviation of the point estimate of 0.64 by [Blundell et al. \(2008\)](#). This result too is driven by the age-dependent persistence of shocks and can be understood in light of [Kaplan and Violante \(2010\)](#)'s finding that a persistent, but not permanent, earnings process can imply the “right” level of insurance against persistent shocks. The richer earnings process implies a significantly lower degree of persistence of the persistent earnings component, particularly at younger ages, compared to the canonical process. Consequently, persistent shocks, particularly at younger ages when assets are low, are more easily self-insured through borrowing and lending.

Our third finding is that our rich earnings process does not improve the fit of the right tail of the wealth distribution with respect to the canonical earnings process.¹³ This is perhaps not so surprising given an established literature, surveyed in [De Nardi and Fella \(2017\)](#), pointing to the fact that accounting for the saving of the rich requires mechanisms—such as a non-homothetic bequest motive, medical-expense risk and entrepreneurship—that go beyond idiosyncratic earnings risk.

Finally, from a normative perspective we find that the welfare costs of earnings risk—as measured by the yearly consumption equivalent—are 2 percentage points lower under the richer earnings process than under the canonical one. The main reason for this finding is again that, while under the canonical process earnings have a permanent, random-walk, component, the richer process implies a lower persistence, particularly in the first part of the working life. As a result, life-cycle risk can be more effectively self-insured under the richer earnings process. Interestingly, the reduction in welfare costs would be even

¹³In [De Nardi, Fella and Paz-Pardo \(2016\)](#) we show that this conclusion still holds if we estimate a similar richer process on synthetically generated W2 data from the parametric processes proposed in [Guvenen et al. \(2016\)](#). It is thus not related to issues of lack-of-oversampling and non-participation by higher income people that are usually associated with most survey data sets.

higher—4 rather than 2 percentage points—in the absence of non-normality and nonlinearities that partly offsets the welfare gains due to age-dependent persistence and innovation variances.

An additional contribution of this Chapter is to propose a simple, simulation-based, method to discretize nonlinear and nonnormal stochastic processes to introduce them in a computational model. Standard discretization methods used in macroeconomics, such as [Tauchen \(1986\)](#) and [Rouwenhorst \(1995\)](#), require the continuous process to be approximated to be linear, typically an AR(1), and, in the case of [Tauchen \(1986\)](#), normal.¹⁴ Our method applies to any, otherwise unrestricted, age-dependent, first-order Markov process. It relies on simulating a panel of individual earnings histories using the continuous process to be approximated and estimating an age-specific, first-order Markov chain on it. This is achieved by discretizing the simulated marginal distribution of earnings at each age—e.g. into percentiles—and by replacing the (heterogeneous) values of earnings in each rank percentile with their median. The associated, age-specific transition matrix is then obtained by computing the proportion of observations transiting from each percentile rank of the earnings distribution at age t to that at age $t + 1$. The result is a non-parametric representation of the process that follows a Markov chain with an age-dependent transition matrix and a fixed number of age-dependent earnings states.

This Chapter is related to the econometric literature on earnings dynamics¹⁵ as well as the macroeconomic literature on the relationship between income and consumption and wealth inequality over the life cycle. [Deaton and Paxson \(1994\)](#) is the seminal empirical contribution. [Storesletten et al. \(2004b\)](#), [Guvenen \(2007\)](#), [Primiceri and Van Rens \(2009\)](#), [Huggett, Ventura and Yaron \(2011\)](#) and [Guvenen and Smith \(2014\)](#) analyze lifetime inequality from the perspective of the standard, incomplete markets model as we do here. Within this literature, many of the consequences of richer earnings processes on consumption, savings and welfare in structural models are still unexplored, with few exceptions. [Castañeda et al. \(2003\)](#) propose an “awesome or superstar” shock to earnings that is unlikely to be observed in the data but that can help to explain the emergence of super-rich people. [Karahan and Ozkan \(2013\)](#) study

¹⁴[Fella, Gallipoli and Pan \(2017\)](#) show how [Tauchen \(1986\)](#) and [Rouwenhorst \(1995\)](#) can be extended to allow for age dependence. Their method still requires linearity though.

¹⁵Besides [Arellano et al. \(2017\)](#) and [Guvenen et al. \(2016\)](#), discussed above, it includes [Geweke and Keane \(2000\)](#), [Lillard and Willis \(1978\)](#), [Bonhomme and Robin \(2009\)](#), [Meghir and Pistaferri \(2004\)](#), [Blundell, Graber and Mogstad \(2015\)](#), [Browning, Ejrmaes and Álvarez \(2010\)](#), and [Altonji, Smith and Vidangos \(2013\)](#). Recent developments are discussed in [Meghir and Pistaferri \(2011\)](#).

the implications of age-dependent persistence and variance of shocks. [McKay \(2017\)](#) finds that taking into account the procyclical nature of negative skewness in earnings growth rates substantially raises the volatility of aggregate consumption growth. [Golosov, Troshkin and Tsyvinski \(2016\)](#) show that higher order moments of earnings shocks are important determinants of optimal redistribution and insurance. [Civale, Díez-Catalán and Fazilet \(2016\)](#) study the implications of skewness and kurtosis for the aggregate capital stock in an economy à la [Aiyagari \(1994\)](#).

The rest of the Chapter is organized as follows. Section [2.2](#) describes the main features of the data on earnings dynamics for both individuals and households. Section [2.3](#) details the methods we use to estimate the canonical and nonlinear earnings processes and their implications. Section [2.5](#) explains the discretization procedure we propose to tractably introduce rich nonlinear earnings dynamics in a quantitative life-cycle model. Section [2.6](#) presents the model and its calibration. Section [2.7](#) discusses the consumption, wealth, and welfare implications of the two earnings processes that we consider, and decomposes the determinants of their differences. Section [2.8](#) concludes. The Appendix discusses key features of the PSID data, details of the estimation and reports a number of robustness checks.

2.2 Earnings Data and their Features

Recent empirical literature has called into question the established view that (log) earnings dynamics are well approximated by a linear model of which the canonical random-walk permanent/transitory model ([Abowd and Card, 1989](#)) with normal innovations is a popular example. Linear models imply that persistence and other second and higher moments are independent of earnings histories. Instead, [Guvenen et al. \(2016\)](#) and [Arellano et al. \(2017\)](#) document that, contrary to the implications of the canonical model, individual pre-tax earnings display both substantial deviations from log-normality and non-linearity.

[Guvenen et al. \(2016\)](#) use confidential Social Security Administration (W2) tax data to establish these facts. The W2 data set has both advantages and disadvantages compared to the PSID data (and household survey data sets more generally). Regarding its advantages, the W2 data set has a large number of observations, is less likely to be contaminated by measurement error, and is not affected by top-coding and differential survey responses. Thus, it could provide better information on the top earners to the extent that they do not respond to

surveys but do pay taxes on all of their earnings. An important disadvantage of the W2 data set is that it is collected at the individual level and lacks the information to identify households and thus to construct household earnings.

The latter is an important shortcoming. In the U.S., the majority of adults are married, 95% of married couples file their income taxes jointly, and taxation of couples and singles is different. Therefore, one needs to know the earnings of both people in a household in order to compute disposable earnings. In this respect household survey data sets that keep track of household structure, like the PSID, have a distinct advantage. This is particularly important if, as we do here, one wants to understand the implications of earnings risk for consumption insurance, which requires taking into account that households and taxes provide insurance against earnings shocks. For such a purpose, *disposable household earnings* is the relevant variable of interest.

The data used in this Chapter are from the Panel Study of Income Dynamics (PSID), 1968-1992. Our sample consists of households who are in the representative core sample, whose head is between 25 and 60 years of age. Given the focus on the implications of earnings risk for consumption insurance, our main variable of interest is *disposable household* labor earnings, although we also discuss the properties of individual pre-tax labor earnings for the purpose of comparison with some closely related work (e.g. [Arellano et al., 2017](#); [Güvenen et al., 2016](#)).

Disposable, household labor earnings are defined as the sum of household labor income and transfers, such as welfare payments, net of taxes and Social Security contributions paid. Appendix B contains a more detailed description of the PSID data we use, our definition of household earnings and how we estimate taxes on labor following [Güvenen and Smith \(2014\)](#).

We adjust our earnings measure for demographic differences across households, since these have no counterpart in the model in Section 2.6. We do so by regressing log earnings on family size. We apply the same regression to the CEX consumption data we use to compute the moments reported in Section 2.7. The residuals from these regressions are interpreted as earnings and consumption per-adult equivalent in the analysis below. For both earnings and consumption, we extract business cycle effects by running a regression of their log levels on year dummies. Finally, we separate the predictable from the stochastic component of earnings by running a regression of our equalized earnings measure on age dummies. We use the residuals to estimate the

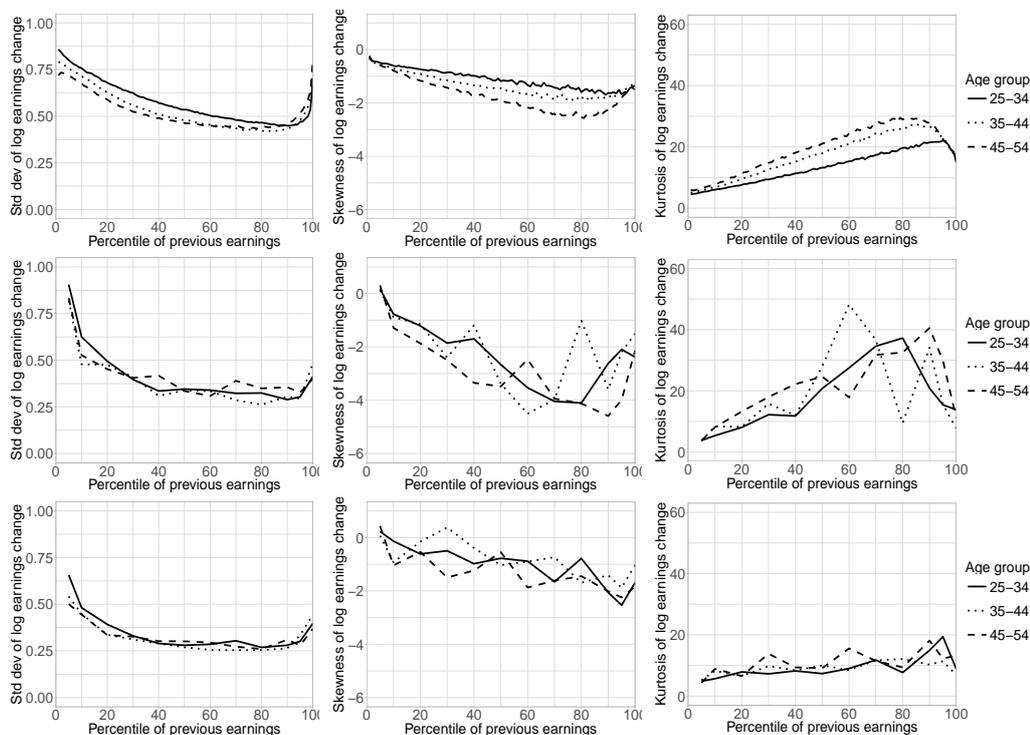


Figure 11: Standard deviation, skewness, and kurtosis of male pre-tax earnings growth in the W2 (top row) and PSID (middle row), and of household after-tax earnings growth in the PSID (bottom row)

stochastic processes for earnings in what follows.

2.2.1 Individual Pre-tax Earnings in the PSID and the W2 Data

We now turn to comparing the properties of *individual pre-tax* earnings data in the PSID with those in the W2 data reported by [Guvenen et al. \(2016\)](#).

The top two rows of Figure 11 compare the second to fourth moments of the W2 data and the PSID. The top row, derived from the moments reported by [Guvenen et al. \(2016\)](#), plots the conditional standard deviation, skewness and kurtosis (measured as the third and fourth standardized moments) of *individual pre-tax* log earnings growth in the W2 data by age and percentile of previous earnings. The middle row of the same figure reports the same moments, by age and decile of previous earnings, for the PSID.¹⁶

Comparing these two sets of figures shows that, overall, the moments in the PSID data are both qualitatively and quantitatively close to those computed from the W2 data. More specifically, the conditional standard deviation of *in-*

¹⁶For comparability with [Guvenen et al. \(2016\)](#), we report moments for households whose head is a male. All moments are very close to those including female household heads. We show moments for 10-year age groups, which we obtain, in the W2 case, by averaging the moments reported by [Guvenen et al. \(2016\)](#) with equal weights for each 5-year age group.

individual pre-tax log earnings growth is U-shaped across all age groups, declining until the 40th percentile and increasing again from the 90th percentile onwards. The increase is more pronounced for the top percentiles in the W2, likely reflecting the coarser partition of the distribution in the PSID data. The most notable difference is the much higher variance at all percentiles above the 20th in the W2 data.

The figures also show that in both datasets *individual pre-tax* log earnings growth has strong negative skewness and very high kurtosis, and that these moments differ by both age and previous earnings. Skewness is more negative for individuals in higher earnings percentiles and for individuals in the 45-54 age group. This indicates that individuals face a larger downward risk in middle age.¹⁷ The comparison of the implications of the two data sets also reveals that, if anything, there is more negative skewness in the PSID data than in the W2 data, except perhaps at the lowest earnings percentiles.

The kurtosis of *individual pre-tax* log earnings growth is hump-shaped by earnings percentile. Even for kurtosis, the maximum value is higher in the PSID, 40, against 30 in the W2 data (compared to 3 for a standard normal distribution).¹⁸

Taken together, these moments provide strong evidence against the standard assumption of a log-normal and linear earnings process for *individual pre-tax* log earnings growth for the PSID data, as well as for the W2 data.

2.2.2 Individual Pre-tax versus Household Disposable Earnings in the PSID

Now that we have shown that *individual pre-tax* earnings growth in both the W2 and PSID data displays remarkably similar deviations from normality and linearity, we turn to contrasting the properties of *individual pre-tax* and *disposable household* earnings in the PSID.

The bottom row in Figure 11 reports the same moments as the first two rows but for *disposable household* log earnings growth (bottom row) in the PSID. Comparing the bottom to the middle row reveals that, as one might have expected, disposable household earnings display lower variance, skewness,

¹⁷Graber and Lise (2015) generate this kind of earnings behavior in the context of a search and matching model with a job ladder.

¹⁸The levels and profiles of skewness and kurtosis of *individual pre-tax* log earnings growth are similar in the two datasets also when looking at report robust measures that exclude outliers (Kelly skewness and Crow-Siddiqui kurtosis). The main difference is a higher level of Crow-Siddiqui kurtosis in the W2 data than in the PSID (see Appendix 2.3).

and kurtosis than pre-tax individual earnings. More specifically, the standard deviation is about 20% smaller at the lower end of the distribution of previous earnings, while skewness and kurtosis are about half as large. Thus, households and taxes perform an important insurance role in buffering individuals from pre-tax earnings changes (as shown by [Blundell, Pistaferri and Saporta-Eksten \(2016\)](#)). This has to be taken into account when considering the economic implications of earnings shocks.

To sum up, the above discussion has shown that, even after taking into account the insurance implied by pooling at the household level and the tax and welfare system, labor earnings display features that contrast with the *age-independence, normality, and linearity* (independence of variance, skewness and kurtosis of previous earnings realizations) implied by the canonical earnings process.

The same is true of another aspect on nonlinearity, nonlinear persistence, that has been documented by [Arellano et al. \(2017\)](#) using *pre-tax* earnings from the PSID. Figure 12 shows how this same feature is prominent also for *disposable* household earnings. It reports earnings persistence as a function of both the previous- and current-earnings rank in our PSID sample. In line with [Arellano et al. \(2017\)](#)'s findings, we also find that earnings persistence is lower (about 0.6) when previous earnings are highest and the current earning shock is lowest and when previous earnings are lowest and the current earning shock is highest (0.4).

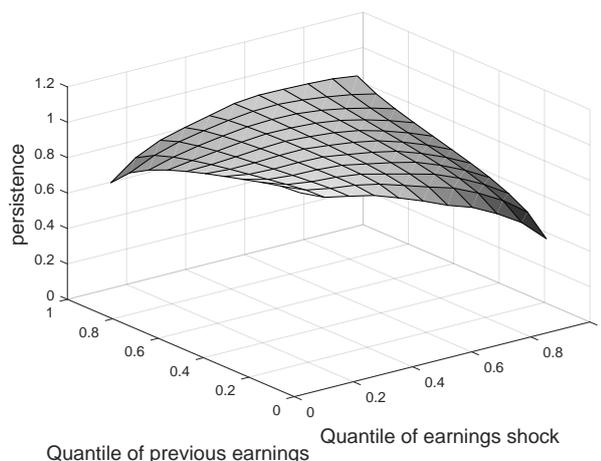


Figure 12: Persistence in log-earnings as a function of previous earnings rank and the rank of the shock received in the current period. PSID data.

2.3 Earnings Processes and their Estimation

2.3.1 Earnings Processes

We start by introducing the canonical linear model of earnings dynamics used in macroeconomics before presenting its nonlinear generalization in [Arellano et al. \(2017\)](#).

Consider a cohort of households indexed by i and denote by $t = 1, \dots, T$ the age of the household head. Let y_{it} denote the logarithm of (residual) disposable household earnings for household i at age t which can be decomposed as

$$y_{it} = \eta_{it} + \epsilon_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T \quad (6)$$

where η and ϵ are assumed to have absolutely continuous distributions. The first component, η_{it} , is assumed to be *persistent* and to follow a first-order Markov process. The second component, ϵ_{it} , is assumed to be *transitory*, have zero mean, be independent over time and of η_{is} for all s .

The *canonical (linear) model* used in macroeconomics is described by

$$\eta_{i,t} = \rho\eta_{i,t-1} + \zeta_{it}, \quad (7)$$

$$\eta_{i1} \stackrel{id}{\sim} N(0, \sigma_{\eta_1}), \quad \zeta_{it} \stackrel{iid}{\sim} N(0, \sigma_{\zeta}), \quad \epsilon_{it} \stackrel{iid}{\sim} N(0, \sigma_{\epsilon}). \quad (8)$$

Thus, the persistent component η_{it} is an autoregressive process of order one with the innovation ζ_{it} independent of $\eta_{i,t-1}$, while the transitory component ϵ_{it} is white noise.

Equations (7)-(8) impose three types of restrictions

1. *Age-independence* (stationarity) of the autoregressive coefficient ρ and of the shock distributions, which imply age-independence of the second and higher moments of the conditional distributions of both the transitory and the persistent component. This is clearly at odds with the strong age-dependence in [Figure 11](#).
2. *Normality* of the shock distributions, which is inconsistent with the negative skewness and high kurtosis discussed above.
3. *Linearity* of the process for the persistent component. Linearity implies:
(a) the additive separability of the right hand side of equation (7) into the conditional expectation—the first addendum—and an innovation ζ_{it} independent of $\eta_{i,t-1}$, and
(b) the linearity of the conditional expectation in

$\eta_{i,t-1}$. Under separability, deviations of η_{it} from its conditional expectation are just a function of the innovation ζ_{it} . As a consequence, all conditional centered second and higher moments are independent of previous realizations of η . This is clearly inconsistent with the dependence of the moments reported in figures 11 and 12 on previous earnings realizations.

The evidence discussed in Section 2.2.1 is what motivates us to consider a more general process that relaxes the above three restrictions while maintaining the first-order Markov assumption for η . The question of how to easily introduce a richer and yet tractable earnings process in a structural model is non-trivial and part of what we propose in this Chapter.

We proceed in two steps. First, we use the quantile-based panel data method proposed by Arellano et al. (2017) to estimate a non-parametric model that allows for age-dependence, non-normality and nonlinearity, and that can be applied in datasets of moderate sample size like the PSID. This step gives us quantile functions for both the two (persistent and transitory) component of earnings. Second, we use the two quantile functions to simulate histories for the two earnings components and proceed to estimate, for each of them, a discrete Markov-chain approximation, which can then be easily introduced in a structural model (this latter step is discussed in detail in Section 2.5).

Let $Q_z(q|\cdot)$, the conditional quantile function for the variable z , denote the q th conditional quantile of z .¹⁹ The process for η can be written in a very general form by replacing equation (7) with

$$\eta_{it} = Q_\eta(v_{it}|\eta_{i,t-1}, t), \quad v_{it} \stackrel{iid}{\sim} U(0, 1), \quad t > 1. \quad (9)$$

Intuitively, the quantile function maps random draws v_{it} from the uniform distribution over $(0, 1)$ (cumulative probabilities) into corresponding random (quantile) draws for η . In the linear case in equation (7) the quantile function specializes to the linearly separable form $Q_\eta(v_{it}|\eta_{i,t-1}, t) = \rho\eta_{i,t-1} + \phi^{-1}(v_{it}; \sigma_\zeta)$, where $\phi^{-1}(v_{it}; \sigma_\zeta)$ is the inverse of the cumulative density function of a normal distribution with zero mean and standard deviations σ_ζ . So, age-independence, normality, and linearity can be seen as restrictions on the quantile function in equation (9).

In particular, one way to understand the role of nonlinearity is in terms of

¹⁹Intuitively, the conditional quantile function is the inverse of the conditional cumulative density function of the variable z mapping from the $(0, 1)$ interval into the support of z . Namely, $z_q = Q_z(q|\cdot)$ satisfies $P[z \leq z_q|\cdot] = q$, where $P[\cdot|\cdot]$ denotes the conditional probability.

a generalized notion of persistence

$$\rho(q|\eta_{i,t-1}, t) = \frac{\partial Q_\eta(q|\eta_{i,t-1}, t)}{\partial \eta_{i,t-1}} \quad (10)$$

which measures the persistence of $\eta_{i,t-1}$ when it is hit by a shock that has rank q . In the canonical model, $\rho(q|\eta_{i,t-1}, t) = \rho$, independently of both the past realization of $\eta_{i,t-1}$ and of the shock rank q . Instead, the general model allows persistence to depend both on the past realization $\eta_{i,t-1}$, but also on the sign and magnitude of the shock realization. Basically, in the nonlinear model shocks are allowed to wipe out the memory of past shocks or, equivalently, the future persistence of a current shock may depend on future shocks.

Of course, a similar unrestricted representation can be used for the transitory component ϵ_{it} and the initial condition η_1 , with the only difference that they are not persistent.

2.4 Estimating the Nonlinear Earnings Process

Following [Arellano et al. \(2017\)](#), we parameterize the quantile functions for the three variables as low order Hermite polynomials

$$Q_\epsilon(q|age_{it}) = \sum_{k=0}^K a_k^\epsilon(q) \psi_k(age_{it}) \quad (11)$$

$$Q_{\eta_1}(q|age_{i1}) = \sum_{k=0}^K a_k^{\eta_1}(q) \psi_k(age_{i1}) \quad (12)$$

$$Q_\eta(q|\eta_{i,t-1}, age_{it}) = \sum_{k=0}^K a_k^\eta(q) \psi_k(\eta_{i,t-1}, age_{it}) \quad (13)$$

where the coefficients $a_k^i(q)$, $i = \epsilon, \eta_1, \eta$, are modelled as piecewise-linear splines in q on a grid $\{q_1 < \dots < q_L\} \in (0, 1)$.²⁰ The intercept coefficients $a_0^i(q)$ for q in $(0, q_1]$ and $[q_L, 1)$ are specified as the quantiles of an exponential distribution with parameters λ_1^i and λ_L^i .

If the two earnings components ϵ_{it} and η_{it} were observable one could compute the polynomial coefficients simply by quantile regression for each point of the quantile grid q_j . To deal with the latent earnings components, the estimation algorithm starts from an initial guess for the coefficients and iterates sequentially between draws from the posterior distribution of the latent persistent

²⁰Following [Arellano et al. \(2017\)](#), we use tensor products of Hermite polynomials of degrees (3,2) in $\eta_{i,t-1}$, and age for $Q_\eta(q|\eta_{i,t-1}, age_{it})$ and second-order polynomials in age for $Q_\epsilon(q|age_{it})$ and $Q_{\eta_1}(q|age_{i1})$.

σ_ϵ^2	$\sigma_{\eta_1}^2$	σ_ζ^2	ρ
0.0620	0.2332	0.0060	1
(0.0020)	(0.0061)	(0.0004)	*

Table 1: Estimates for the canonical earnings process (standard errors in parentheses)

components of earnings and quantile regression estimation until convergence of the sequence of coefficient estimates.

Reported standard errors are computed by bootstrapping. In particular, we sample with replacement pairs of observations y_t, y_{t+1} from our PSID sample and then run the estimation algorithm for a large number of those samples.

2.4.1 Estimating the Canonical Linear Earnings Process

We estimate the canonical process for residual earnings in equations (6)-(8) using GMM, where the target moments are variance and autocovariance age profiles in the data.²¹ The associated standard errors are obtained by bootstrapping. Table 1 reports our results. As common in the literature, we find that the persistent component has a unit root.²² For this reason, though we have allowed for individual fixed effects at the estimation stage, their variance cannot be identified separately from the variance of the initial condition $\sigma_{\eta_1}^2$ and we have normalized it to zero.

2.4.2 Comparing the Implications of the Nonlinear and Canonical Earnings Processes

To understand the economic implications of the nonlinear and canonical earnings processes, it is useful to compare their implications in terms of (a) age-dependence of second moments; (b) non-normality; (c) nonlinearity.

Starting from the **age-dependence** of second moments, the top row of Figure 13 plots the age profile of the standard deviations of the shocks to the persistent and transitory components of earnings. Both are age-independent by construction in the canonical process. The standard deviation of shocks to the persistent component is substantially higher for the nonlinear process and

²¹Appendix 2.1.4.2 provides more information about our estimation method.

²²The unrestricted GMM estimation returns an estimate of $\rho = 1.01$. Given that there has been little exploration of the explosive case in the literature, and that we want our canonical process to be standard, the above estimates are obtained under the restriction $\rho \leq 1$. The resulting estimate is at the bound.

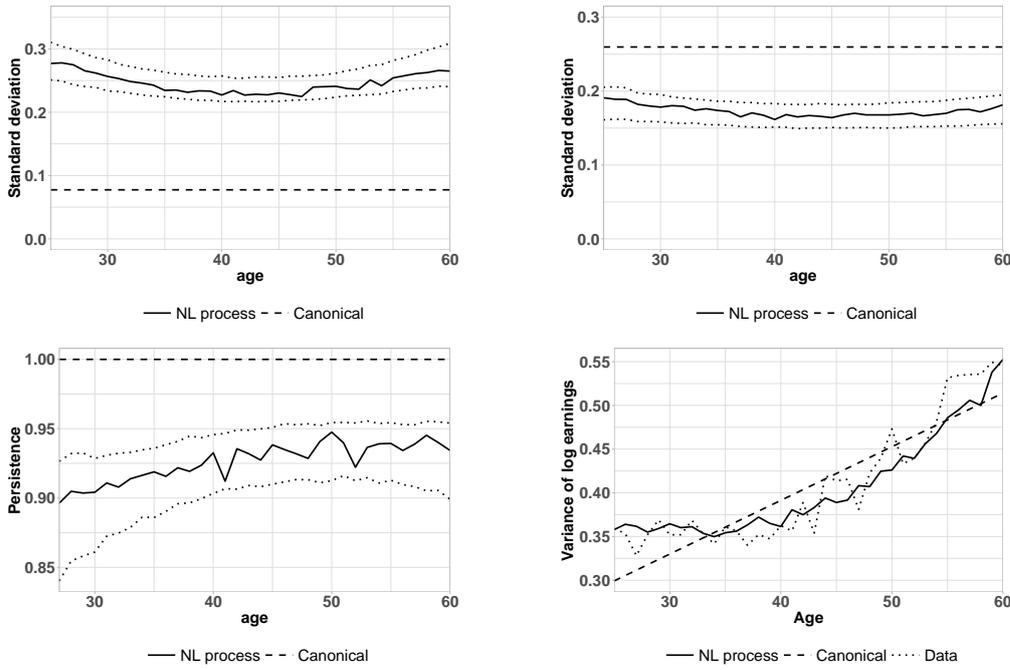


Figure 13: Age dependence of second moments: nonlinear vs canonical process. Top left, standard deviation of the innovation to the persistent component. Top right, standard deviation of the transitory shock. Bottom left, autocorrelation of the persistent component. Bottom right, cross-sectional variance of log earnings. The dotted bands represent bootstrapped 95% confidence intervals.

follows a U-shaped pattern by age. In contrast, the standard deviation of the transitory component of the nonlinear process displays little age variation and is lower in the nonlinear than in the canonical model. The bottom left panel of Figure 13 reports the age-profile of the first-order autocorrelation of the persistent earnings component for the two processes. In the nonlinear earnings process it is lower than in the canonical case for all ages, but it does increase between age 25 and 45. Given these differences, it is not surprising that the nonlinear process provides a substantially better fit of the age profile of the cross-sectional earnings dispersion, which we display in the bottom right panel of Figure 13²³. More specifically, the canonical earnings process cannot capture the convex shape of the cross-sectional variance of earnings by age while the nonlinear process provides an extremely close fit, thanks to the combination of increasing persistence and declining variance of the persistent component over the ages 25 to 45. It is also apparent that the canonical model requires a low variance of the persistent shocks relative to the transitory ones to match the relatively low rate of growth of the cross-sectional variance of earnings over the life-cycle. Figure 14 displays more evidence on age-dependence, which also

²³See Appendix 2.2 for details on the computation of this variance.

manifests itself in the skewness and kurtosis of the shocks.

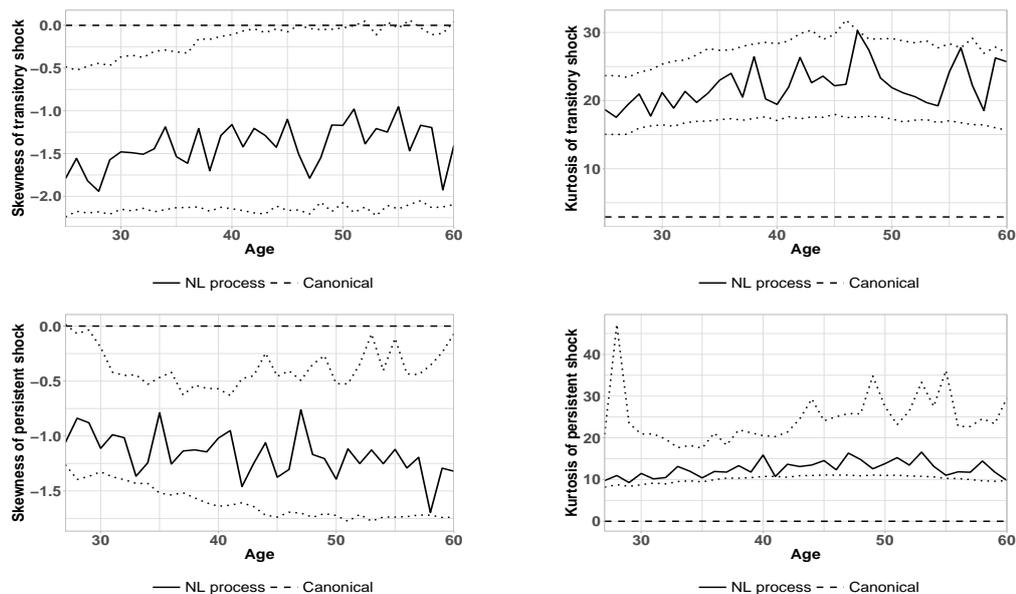


Figure 14: Skewness and kurtosis (by age) of the innovations to (a) the transitory component of earnings (top) and (b) the persistent component of earnings (bottom). The dotted bands represent bootstrapped 95% confidence intervals.

Turning to **non-normality**, Figure 14 reports skewness and kurtosis for the innovation to the transitory (top row) and persistent component of earnings (bottom row) by age and highlights that the earnings data display deviations from normality (the dashed line). They also highlight that non-normality, in particular kurtosis, is higher for transitory than persistent shocks.

Turning to **nonlinearity**, Figure 15 plots the standard deviation of shocks to the innovation to the persistent component of earnings by previous earnings, while the right panel plots the persistence measure in equation (10)—namely the correlation between the percentile of η_{t-1} and of the innovation to it—averaged by age. In the right panel of this figure, we do not plot the persistence of the canonical model (which is constant at 1) for the sake of readability. These two panels clearly illustrate that the constant variance and persistence implied by the canonical process are strongly at odds with the highly nonlinear patterns in Figure 15 and the features of the observed data.

2.5 The Discretized Nonlinear Earnings Process

To use the estimated process (6)-(8) in the life-cycle model, we discretize it using an age-dependent Markov chain.

We start by simulating a large set of histories for the persistent and tran-

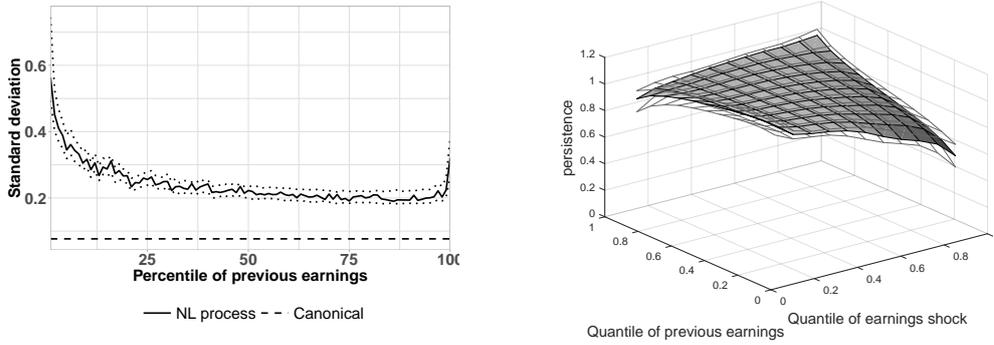


Figure 15: Standard deviations of persistent shocks by previous earnings (left) and nonlinear persistence of the persistent component, by quantile of previous earnings and quantile of shock received in the current period (right). The dotted bands and transparent mesh represent bootstrapped 95% confidence intervals.

sitory component of earnings. For each component in the simulated sample, we estimate a Markov chain of order one, with age-dependent state space $Z^t = \{\bar{z}_1, \dots, \bar{z}_N\}$, $t = 1, \dots, T$ and an age dependent transition matrices Π^t , of size $(N \times N)$. That is, we assume that the dimension N of the state space is constant across ages but we allow the set of states and the transition matrices to be age-dependent.

We determine the points of the state-space and the transition matrices at each age in the following way.

1. At each age, we order the realizations of each component by their size and we group them into N bins. In our main specification the choice of bins reflects the tradeoff between approximating reasonably well the rich earnings dynamics and while maintaining a reasonably low number of bins. For the persistent component we use 18 bins with each bin representing deciles, with the exception of the top and bottom deciles, that we split in 5. Thus, bins 1 to 5 and 14 to 18 include 2% of the agents at any given age, while bins $n = 6, \dots, 13$ include 10% of the agents at any given age. For the transitory component we use 8 bins, with bins 1 to 2 and 7 to 8 including 2.5% of agents, 3 and 6 including 5% and 2 and 3 each having 40% of agents.²⁴
2. The points of the state space at each age t are chosen so that point z_t^n is the median in bin n at age t . Kennan (2006) proves that setting the gridpoint at the median of the bin (in the specific case of equally-sized

²⁴Bin sizes were chosen heuristically on the basis of quality of fit of the moments given a reasonable bin number. Given non-normality, this required a finer partition for the tails of the distributions. The Online Appendix D presents robustness results concerning this discretization.

bins) and attributing a weight of $1/N$ to each of the N bins constitutes the best discrete approximation of an arbitrary distribution.

3. The initial distribution at model age 0 is the empirical distribution at the first age we consider.
4. The elements π_{mn}^t of the transition matrix Π^t between age t and $t + 1$ are the proportion of individuals in bin m at age t that are in bin n at age $t + 1$.

Allowing for an age-dependent Markov chain allows to capture the non-constancy of moments of the earnings distribution over the life-cycle. The flexible form of the transition matrix allows to capture nonlinearities as a function of current earnings. The use of this kind of transition matrices is well established in the literature. [Krueger and Perri \(2003\)](#) use them to study the welfare consequences of an increase in earnings inequality. Studies of income mobility (e.g. [Jäntti and Jenkins \(2015\)](#)) and consumption mobility (e.g. [Jappelli and Pistaferri \(2006\)](#)) rely on them to analyze intra- or inter- generational mobility across relative rankings in the distributions. In this study, instead, we are interested in capturing movements across earnings levels.

2.6 The Model

The model is a partial-equilibrium, life-cycle, incomplete-markets model in the tradition of [Bewley \(1977\)](#). There is no aggregate uncertainty.

2.6.1 Demographics

Each year, a positive measure of agents is born. People start life as workers and work until retirement at age T^{ret} . The population grows at rate n .

An agent of age t faces a positive probability of dying ($1 - s_t$) by the end of the period, where s_t denotes the one-period survival probability. Agents die with probability one by age T .

2.6.2 Preferences and Endowments

Preferences are time separable, with a constant discount factor β . The intra-period utility is CRRA: $u(c_t) = c_t^{1-\sigma}/(1 - \sigma)$.

Agents are endowed with one indivisible unit of labor which they supply inelastically at zero disutility. Their earnings are subject to random shocks and follow the process described by equations (6)-(8).

2.6.3 Markets and the Government

Asset markets are incomplete. Agents can borrow up to an exogenous borrowing limit \underline{a} and can only invest in a risk-free asset at an exogenous rate of return r . There are no annuity markets to insure against mortality risk. As a result, there is a positive flow of accidental bequests in each period. We assume these are lost to the economy and thus are not received by any individual or the government.

Retired individuals receive an after-tax pension p from the government until they die. The pension is a function of the last earnings realization.

2.6.4 The Household's Problem

In any given period, a agent of age t chooses consumption c and risk-free asset holdings for the next period a' , as a function of the relevant state vector. The optimal decision rules for consumption and savings solve the dynamic programming problems described below.

- (i) Agents of working age $t < T_{ret}$ solve the recursive problem

$$\begin{aligned} V(t, z, \eta) &= \max_{c, a'} \left\{ u(c) + \beta s_t E_t V(t+1, z', \eta') \right\} & (14) \\ \text{s.t. } a' &= z - c, \quad a' \geq \underline{a}, \\ z &= (1+r)a + \eta + \epsilon, \end{aligned}$$

where z denotes total cash at hand.²⁵

- (ii) From the retirement age T^{ret} to the terminal age T agents no longer work and live off their pension p and accumulated wealth. Their value function satisfies:

$$\begin{aligned} W(t, z, p) &= \max_{c, a'} \left\{ u(c) + s_t \beta W(t+1, z', p) \right\} & (15) \\ \text{s.t. } a' &= z - c, \quad a' \geq \underline{a}, \\ z &= (1+r)a + p. \end{aligned}$$

The agent's pension p enters the state vector because it is a function of the agent's earnings pre-retirement. The terminal value function $W(T, a, p)$ is equal to zero (agents do not derive utility from bequests).

²⁵The choice of state vector avoids separately keeping track of the transitory component of earnings ϵ which is independently distributed over time.

2.6.5 The Model Calibration

The model period is one year. Agents enter the labor market at age 25. The retirement and terminal ages are $T^{ret} = 60$ and $T = 85$. The population growth rate n is set to 1.2% per year. The survival probabilities s_t are from [Bell, Wade and Goss \(1992\)](#).

The coefficient of relative risk aversion is set to 2, a standard value. The risk-free rate is 4% and the discount factor β is calibrated to match a wealth to income ratio of 3.1. It equals 0.957 under the canonical earnings process and 0.939 under the nonlinear one.

We set the exogenous borrowing limit \underline{a} to 12% of average disposable household earnings. This represents the average credit card limit in the SCF in 1989 and 1992, the two years within our sample period for which that information is available in the SCF.

As described in [Section 2.2](#), our earnings processes are based on disposable earnings, hence we do not explicitly include taxation in the model.²⁶ In both cases, we impose the same average income profile, which we estimate from our PSID sample.

We discretize the two earnings processes as follows. In the case of the canonical earnings process, whose estimates we report in [Table 1](#), we discretize the persistent component using the modified version of the Rouwenhorst method for non-stationary processes proposed by [Fella et al. \(2017\)](#). In the case of the nonlinear earnings process, we apply the procedure described in [Section 2.5](#). For both cases, we use 18 gridpoints for the persistent component and 8 for the transitory component. In the Online Appendix we show that our results are robust to increasing the number of gridpoints and alternative discretization procedures.

The Social Security benefit p is a function of the last realization of disposable earnings $y_{ret} = \eta_{ret} + \epsilon_{ret}$, which follows a fixed schedule of replacement rates. Namely, there is a 90% replacement rate for the part of y_{ret} below 18% of average earnings, of 32% for the fraction between 18% and 110% of average earnings, and 15 percent for the remainder. All benefits are then (very slightly) scaled up proportionately so that a worker that makes average earnings is entitled to a 45% replacement rate. The function is meant to mimic the US system and is based on [Kaplan and Violante \(2010\)](#).

²⁶[Appendix B](#) provides more details about the earnings definition.

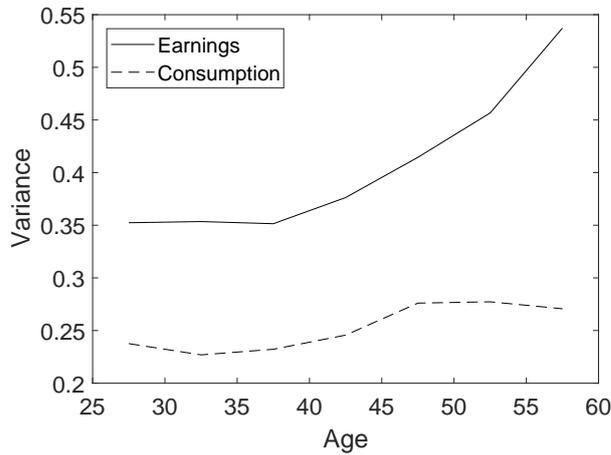


Figure 16: Cross-sectional variance of log earnings and log consumption by age. See Appendix 2.2 for details on their computation.

2.7 Consumption, Wealth, and Welfare Implications

This section studies the model’s implications under the canonical and nonlinear earnings processes and compares them to U.S. consumption data. To do so, we first analyze the growth in consumption dispersion over the working life and then turn to measuring self-insurance insurance as proposed by [Blundell et al. \(2008\)](#). Finally, we compare the implications of these earnings processes for wealth inequality and welfare.

2.7.1 Consumption Inequality over the Working Life

We start by studying the evolution of cross-sectional consumption dispersion over the life cycle. Following [Deaton and Paxson \(1994\)](#) and [Storesletten, Telmer and Yaron \(2004a\)](#), it is common to interpret its growth, relative to the growth of cross-sectional earnings dispersion, as a measure of risk sharing. For reference, Figure 16 plots the cross-sectional dispersion of consumption and earnings over the life cycle. The (solid) earnings variance profile is from our PSID sample, while the dashed line plots the variance profile of nondurable consumption is from the CEX 1980-2007 as in [Heathcote, Perri and Violante \(2010\)](#). Given the relatively small sample size, we group observations in 5-year age groups. As it is well known, both earnings and consumption inequality increase over the working life, but earnings inequality increases substantially faster than consumption inequality from about age 40.

Because the increase in consumption inequality over the working life is informative about people’s ability to insure against earnings risk, it provides a useful benchmark against which to assess the ability of the model to capture

the degree of insurability of earning shocks in the data. Figure 17 reports the age profile of cross-sectional consumption dispersion for both the CEX data and the model under, respectively, the canonical and nonlinear earnings processes.²⁷ The series are normalized so that each starts at zero at age 27, which is the midpoint of the first 5-year age group (25–29). The canonical earnings process fails to match both the overall growth and the shape of the profile of consumption dispersion. Its overall growth rate is more than double that in the data and its profile is monotonic and roughly linear. Conversely, in the data, the age profile of consumption is significantly convex between age 25 and 47. The nonlinear process, instead, matches well both the overall growth in consumption dispersion and its convexity in the first part of the life cycle. The one part that it misses is the flattening out after age 47.

The finding that the estimated richer earnings processes implies a profile of consumption dispersion in line with the data is remarkable. Standard models with linear earnings processes (see Storesletten et al. (2004b)) generate a profile similar to the one implied by the canonical earnings process in Figure 17, and thus overstate the rate of growth of consumption dispersion, unless the process for earnings has an idiosyncratic deterministic time trend, or Heterogeneous Income Profile (Guisar, 2007; Primiceri and Van Rens, 2009). Intuitively, heterogeneity in individual, life-cycle trend growth generates a substantially smaller rise in consumption dispersion because the individual-specific trend growth is known to consumers but not to the econometrician. Huggett et al. (2011) show that heterogeneity in earnings growth rates can be also generated by the endogenous response of human-capital investment over the life cycle to heterogeneity in initial human capital levels.

Our findings suggest a novel explanation: the age profile of cross-sectional consumption dispersion can be generated by the response of saving to the richer earnings dynamics that we consider, without resorting to heterogeneity in income profiles. It should also be noted that allowing for heterogeneity in income profiles cannot generate (cfr. Guisar, 2007; Primiceri and Van Rens, 2009) the strong convexity that characterizes the consumption data (dotted line in Figure 17).

As we have discussed in Section 2.3.1, our rich earnings process deviates from the canonical linear process along three main dimensions: (1) age-dependence,

²⁷We perform this comparison recalibrating beta so as to keep the wealth to income ratio constant across earnings processes.

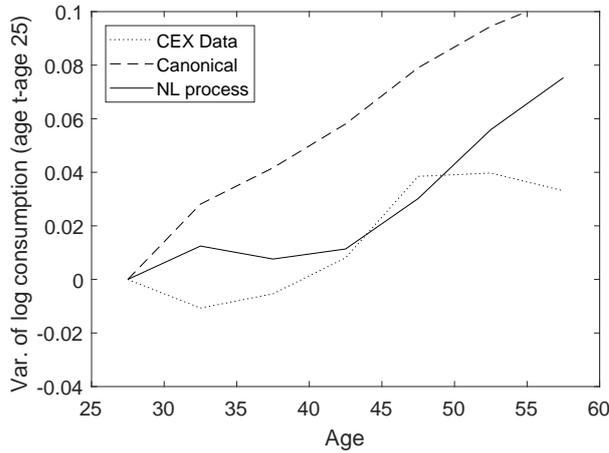


Figure 17: Growth in the cross-sectional variance of log consumption, data and implications of two earnings processes.

(2) non-normality, and (3) nonlinearity. To understand the contribution of each of these factors to the growth of consumption dispersion over the life cycle, we conduct a series of counterfactual experiments, simulating the model under progressively richer stochastic processes for earnings.

We start by restricting the functional form of the earnings process to be the sum of an AR(1) plus a white noise component, as in the canonical process, but allowing for both age-dependent persistence and variance of shocks (as in [Karahan and Ozkan \(2013\)](#)), as well as non-normality of their distributions. Compared to the fully general nonlinear earnings process, this one imposes linearity in $\eta_{i,t-1}$; namely, that persistence and other second and higher conditional moments are independent of $\eta_{i,t-1}$. We estimate this process on our PSID data, following the procedure described in Section 2.3.1 for the nonlinear process, but restricting the quantile function for the persistent component in equation (9) to be linear in its past value.

To further disentangle the effect of the age dependence of persistence and variance from that of non-normality, we perform two set of simulations using the restricted estimates that we have just described. In the first one, we simulate earnings using the estimated persistence and variances but drawing shocks from a normal distribution. In the second experiment, we simulate earnings using the estimated distribution (i.e. quantile function), that also allows for non-normality. We discretize each of the resulting processes using the method in Section 2.5. The recalibrated value of the discount factor equals 0.939 in the economy with normal shocks and 0.940 in the other one.

Figure 18 plots the cross-sectional variance profiles reported in Figure 17,

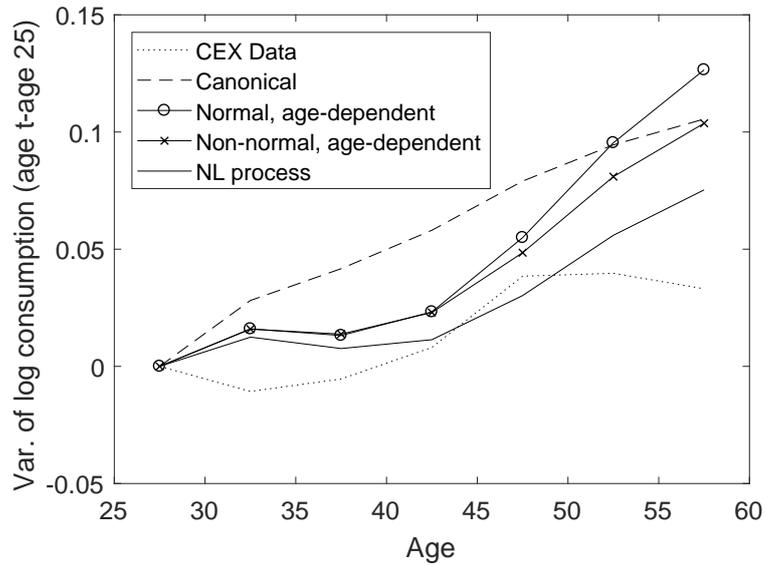


Figure 18: Growth in the cross-sectional variance of log consumption. Contribution of (1) age-dependence, (2) non-normality and (3) nonlinearity to the growth of the cross-sectional variance of log consumption.

with the addition of the two profiles implied by (a) only age-dependence and (b) age-dependence together with non-normality.

The solid line marked with circles in Figure 18 corresponds to the case of an age-dependent linear process with normal innovations. Compared to the canonical case, allowing for age dependence substantially improves the fit of consumption dispersion in the first part of the life cycle, but counterfactually implies an even larger growth rate of consumption dispersion from age 43 onwards. The net effect for the age-dependent earnings process is an overall rate of growth in consumption dispersion between ages 25 and 60 that is three percentage points higher than in the canonical case.

The intuition behind the above finding is the following. Allowing for age-dependence implies that the estimated process for earnings matches the age-profile of the cross-sectional earnings variance in the bottom right panel of Figure 13; namely, relatively flat until age 43 but growing at a rate substantially above its working-life average afterwards. The forces underpinning this pattern are: (a) the U-shaped profile of the variance of the persistent component of earnings; and (b) a persistence below one that increases until age 45 but flattens out afterwards (see Figure 13). Compared to the canonical process with a unit root and constant shock variance, the interaction of these two forces implies that self-insurance through precautionary saving is more effective and, as a consequence, the growth in consumption dispersion is lower until middle age. In the second half of the working life, though, the increase in the variance of

the persistent earnings shocks reduces the ability to self-insure and results in a substantial increase in consumption dispersion. This is confirmed by comparing the age profile of average wealth reported in the left panel of Figure 19 under the canonical (dashed curve) and age-dependent earnings process with normal shocks (solid curve with circles). Though the aggregate wealth-to-earnings ratio is the same in the two economies, average saving is higher before and lower after age 50 in the economy with age-dependent earnings process.

We now turn to the linear process with the same (age-dependent) first and second moments as above but with non-normal innovations. The solid line marked with crosses in Figure 18 plots the associated age profile of variance. Compared to the normal case, the rate of growth of the consumption variance is everywhere lower. The difference is particularly pronounced towards the end of the working life. To understand the mechanism at work, it is important to understand the impact of negative skewness and kurtosis on precautionary saving and the wealth distribution. [Civale et al. \(2016\)](#) study the issue in an Aiyagari economy and show that, everything else equal, negative skewness reduces both the cross-sectional mean and dispersion of wealth while kurtosis increases both.

The effect of higher kurtosis is in line with intuition. By increasing the probability of tail events higher kurtosis increases precautionary saving for all agents and therefore the mean and variance of the wealth distribution. The effect of negative skewness, though, is less intuitive. Basically, for a distribution to have higher negative skewness keeping the other moments constant, some probability mass has to move towards the top of the distribution. Wealthy agents are not sensitive to left skewness but, confronted with a higher probability of positive shocks, save less. Conversely, agents who are close to the borrowing constraint are more sensitive to skewness than to the higher probability of positive shocks and save more. In the aggregate, the response of wealth-rich agents dominates that of the wealth-poor and average wealth falls. More intuitively, so does the variance of wealth holdings. Comparing the lines marked with circles and crosses in Figure 19 reveals that, in our model, the net effect of negative skewness and kurtosis hardly affects the life-cycle profile of average wealth (left panel), but substantially reduces the rate of growth of the variance of wealth holdings (right panel) compared to the case with normal shocks. This fall in wealth dispersion accounts for the fall in consumption dispersion in Figure 18 when skewness and kurtosis of shocks are introduced.

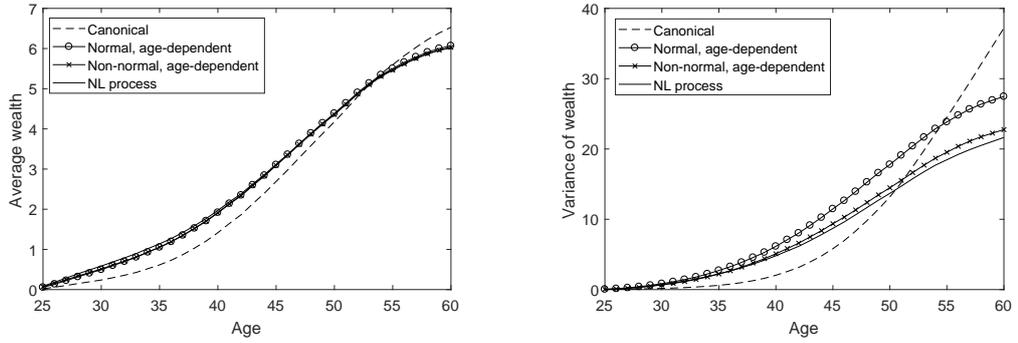


Figure 19: Cross-sectional average wealth (left) and variance of wealth holdings (right), by age and earnings process.

Finally, comparing the line marked with crosses and the solid line in Figure 18 shows that allowing for nonlinearity brings the overall fit of life-cycle inequality substantially closer to the data, compared to all of the other earnings processes considered. Figure 20 provides some insight into the mechanism associated with the nonlinearity in earnings. It plots the persistence (averaged over age) of the persistent earnings component by previous earnings and current shock for both the age-dependent non-normal case (light surface) and the nonlinear one (dark surface). By assumption persistence is constant in the former case. For individuals with previous earnings realizations below the median, negative shocks (below the median) increase persistence relative to the linear case, while positive shocks reduce it. This implies that good shocks partially wipe out the memory of previous bad realizations while bad shocks reinforce it. The average persistence of a bad previous realization is hardly affected but, since the precautionary motive implies that saving responds more to downside than upside risk, individuals with bad earnings realizations save more in the nonlinear case. The nonlinearity is also present, though much less pronounced, for individuals with previous earnings realizations in the top two deciles. For these, shocks below the second decile reduce persistence, while shocks above it increase it, relative to the linear case. The increase in average persistence, relative to the linear case, tends to reduce the saving response. On the other hand, the fact that bad shocks reduce earnings more than linearly (reduce persistence) increases precautionary saving. Overall, saving increases for agents with adverse realizations of previous persistent earnings, while it falls for individuals with good previous earnings realization. This is reflected in the flattening in the age-profile of the variance of wealth holdings in the right panel of Figure 19. This fall in wealth dispersion accounts for the further fall in the growth of

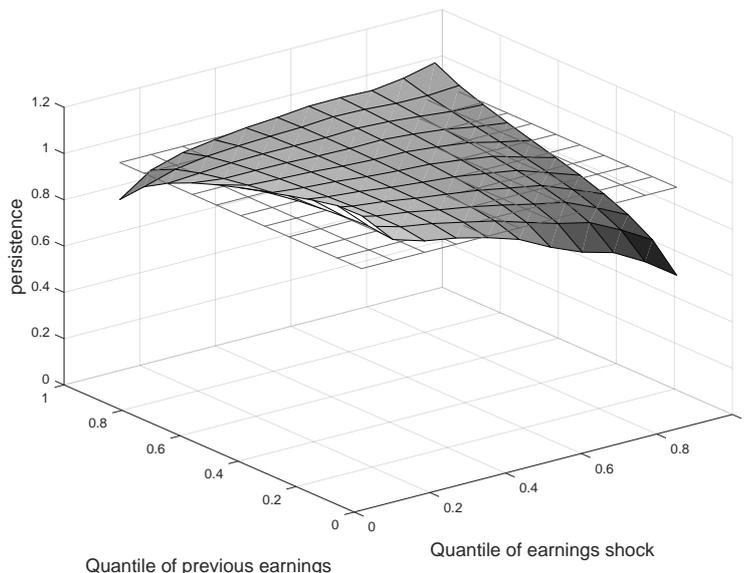


Figure 20: Persistence of the persistent component by quantile of previous earnings and current shock: NL (dark surface) vs non-normal age-dependent process (transparent surface).

consumption dispersion over the life cycle which brings it much closer to the data, particularly for ages up to 50 (see Figure 18).

None of our earnings processes captures the flattening out in the variance of consumption that we measure after age 47 because the variance of earnings in the data keeps increasing. Our structural model misses two aspects of the data that could be important in this regard. The first one is early retirement. For retirees, income is mainly composed of Social Security payments and does not vary much. Thus, consumption is no longer exposed to earnings fluctuations and medical expense risk is not very high until well into retirement age, as shown by [De Nardi, French and Jones \(2010\)](#). The second one is the role of durables and housing, that become substantial by that age and might affect both measured consumption (we only look at nondurable consumption) and one's ability to self-insure.

2.7.2 Measuring Self-Insurance against Earnings Shocks

An alternative, and possibly more intuitive, measure of self-insurance is related to the extent of pass through from shocks to disposable earnings onto consumption. [Blundell et al. \(2008\)](#) propose estimating consumption insurance coefficients on persistent and transitory earning shocks by positing the following equation

$$\Delta c_{it} = (1 - \psi^p)\nu_{it} + (1 - \psi^{tr})\epsilon_{it} + \xi_{it}, \quad (16)$$

where $\nu_{it} = \eta_{it} - E[\eta_{it}|t, \eta_{i,t-1}]$ denotes the innovation to the persistent component of earnings and ϵ_{it} the transitory component. The insurance coefficients with respect to persistent (ψ^p) and transitory (ψ^{tr}) shocks

$$\psi^p = 1 - \frac{\text{cov}(\Delta c_{it}, \nu_{it})}{\text{var}(\nu_{it})}, \quad \psi^{tr} = 1 - \frac{\text{cov}(\Delta c_{it}, \epsilon_{it})}{\text{var}(\epsilon_{it})} \quad (17)$$

capture the fraction of the variance of either type of shock that does not transmute into movements in consumption. Similarly, one can compute age-specific insurance coefficients ψ_t^p, ψ_t^{tr} where moments are computed only over agents of age t .

To compute the insurance coefficients implied by our model, we simulate a panel of working lives under both the benchmark and nonlinear processes and compute the associated consumption c_{it} and insurance coefficients in equation (17) on the simulated data.

Computing the coefficients in equation (17) within the model is straightforward since the shocks are observable. In contrast, estimating them from the data requires identifying the two types of earning shocks at the individual level. [Blundell et al. \(2008\)](#) propose an identification strategy under the assumption that earnings follow the canonical linear process (6)-(8). The estimators for the insurance coefficients based on the BPP methodology are given by

$$\psi_{BPP}^p = 1 - \frac{\text{cov}(\Delta c_{it}, y_{i,t+1} - y_{i,t-2})}{\text{cov}(\Delta y_{it}, y_{i,t+1} - y_{i,t-2})}, \quad \psi_{BPP}^{tr} = 1 - \frac{\text{cov}(\Delta c_{it}, \Delta y_{i,t+1})}{\text{cov}(\Delta y_{i,t}, \Delta y_{i,t+1})}. \quad (18)$$

As pointed out by [Kaplan and Violante \(2010\)](#), comparing the coefficients in equation (18) estimated within the model to the estimates in [Blundell et al. \(2008\)](#) conveys information on the degree of shock insurability in the model relative to the data.

The coefficients in equation (18), though, may provide biased estimates of the true coefficients in equation (17) to the extent that the identification assumption on which they are based is violated. The assumption can be violated for two reasons. First, if earnings do not follow the canonical linear process in equation (6)-(8). This is obviously true in the more flexible cases we consider. Second, as pointed out by [Kaplan and Violante \(2010\)](#), even if earnings follow a canonical linear process the ψ_{BPP}^p estimator may be biased whenever consumption does not equal permanent income, as is the case in the presence of a precautionary saving motive.²⁸ For this reason, we compute both types of

²⁸Formally, the bias is present whenever present consumption responds to past persistent income

Process/Coefficients	ψ_{BPP}^p	ψ_{BPP}^{tr}	ψ^p	ψ^{tr}
Data: BPP (2008)				
Canonical (S.E. in parenthesis)	0.36	0.95	–	–
	(0.09)	(0.04)		
Model				
Canonical	0.13	0.89	0.31	0.92
Nonlinear process	0.43	0.82	0.46	0.91
Normal, age-dependent	0.42	0.83	0.46	0.88
Non-normal, age-dependent	0.42	0.83	0.46	0.88

Table 2: Insurance coefficients

coefficients. Table 2 reports their values under the alternative income processes.

Columns 1 and 2 in Table 2 report the coefficients in equation (18). As a reference, the first row reports the estimates by [Blundell et al. \(2008\)](#)—respectively 0.36 for permanent and 0.95 for transitory shocks—on the PSID using similar data to ours.²⁹ The corresponding values for the model, when earnings follow the canonical earnings process, are 0.13 and 0.89, which confirms the finding by [Kaplan and Violante \(2010\)](#) that the extent of self-insurance of permanent earnings shocks implied by the model is substantially lower than the degree of insurance in the data. On the other hand, the estimates for the model with a nonlinear earnings process imply an insurance coefficient for persistent shocks of 0.43 which is substantially more in line with, and even marginally larger than, the BPP estimate in the first row.

In the case of a canonical income process, the “excess smoothness” ([Campbell and Deaton, 1989](#)) of the consumption response, relative to the predictions of the permanent income model, can be explained on the basis of efficient risk exchange under limited commitment and private information resulting in partial insurance of persistent earnings shocks (e.g. [Attanasio and Pavoni, 2011](#)). Our findings imply that an alternative interpretation is that the observed degree of smoothness is consistent with only self-insurance through risk-free borrowing and lending if earnings follow the estimated nonlinear dynamics.

From a qualitative perspective, this result is very much in line with our

changes, which implies that $\text{cov}(\Delta c_{it}, y_{i,t+1} - y_{i,t-2})$ is a biased estimator of $\text{cov}(\Delta c_{it}, \nu_{it})$. [Kaplan and Violante \(2010\)](#) show that this is indeed the case in a life-cycle model similar to ours with a canonical earnings process and occasionally-binding borrowing constraints.

²⁹[Blundell et al. \(2008\)](#) conduct their analysis using disposable household earnings for continuously married coupled headed by a male head.

findings in Section 2.7 that agents are more able to self-insure against income fluctuations when earnings follow the nonlinear process than in the canonical case. Interestingly, our finding that allowing for a richer earnings process implies a substantially different estimate of the insurance coefficient for persistent shocks is confined to *disposable* household earnings. Using the same earnings process we use here, [Arellano et al. \(2017\)](#) estimate an average insurance coefficient for persistent shocks to *pre-tax* household earnings between 0.6 and 0.7 which is in line with an estimate of 0.69 in [Blundell et al. \(2008\)](#) under the identifying assumption that earnings follow the canonical process. As discussed in [Blundell et al. \(2008\)](#), the nearly double magnitude of the insurance coefficients with respect to *pre-tax* rather than *disposable* earnings is due to the of insurance implied by the tax and transfer system.

Turning to the insurance coefficient for transitory shocks in column 2, it may seem surprising that it is higher under the canonical than under the nonlinear earnings process. As pointed out in [Kaplan and Violante \(2010\)](#), though, the intuition is that the increased insurability of persistent shocks induces households to shift the use of savings from the smoothing of transitory shocks to the smoothing of persistent shocks.

Columns 3 and 4 in Table 2 report the estimates of the *true* insurance coefficients in equation (17) within the model. Comparing them to the BPP estimates in columns 1 and 2 reveals that the downward bias of the insurance coefficient for persistent shocks implied by the BPP procedure is sizeable (0.13 against 0.31) in the case of the canonical income process but small (0.43 against 0.46) for the nonlinear process. The intuition is that, as pointed out by [Kaplan and Violante \(2010\)](#), the bias is exacerbated in an economy in which the borrowing constraint is occasionally binding. As discussed above, when earnings follow the nonlinear process shocks are more insurable, and precautionary saving larger. For this reason, the economy spends less time close to the borrowing constraint and the bias is lower.

Finally, the last two lines reports the same coefficient for the case with age dependence and normal shocks and the one that also allows for non-normality. Comparing the three set of estimates reveals that the feature that drives the better match of the insurance coefficient for persistent shock estimated by BPP is essentially the age dependence of the earnings process. This is consistent with the finding in [Karahan and Ozkan \(2013\)](#) that the (true) insurance coefficient for persistent shocks in a life-cycle economy with an age-dependent earnings

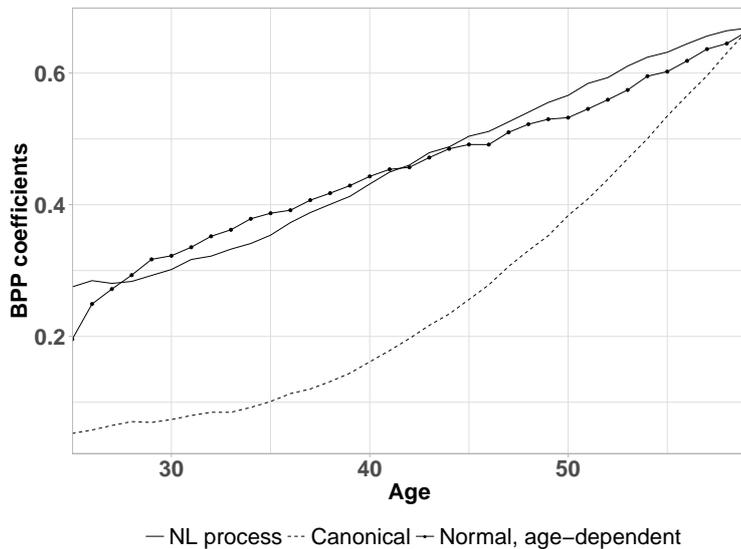


Figure 21: Partial insurance coefficients on persistent shocks, ψ_t^p , by age

process with normal shocks is 0.38.³⁰

While Table 2 reports the average insurance coefficients, Figure 21 plots the true insurance coefficient for persistent shocks ψ_t^p at each age. The coefficients are increasing with age, as: (a) wealth is accumulated; and (b) the fall in the residual working life reduces the effective shock persistence. The degree of insurability at all ages but the last working age is substantially higher under the nonlinear earnings process than under the canonical one. For the same reason, the age profile of the coefficients is substantially flatter in the former case. In line with the discussion above, most of the difference is due to the age-dependence of earnings. It is only from age 40 onwards that the coefficients are marginally higher under the nonlinear process than under the age-dependent earnings process with normal shocks.

2.7.3 Wealth

Table 3 compares the implied wealth distribution of the canonical and nonlinear earnings processes with data from the U.S 1989 SCF (Kuhn and Ríos-Rull, 2015).

As known in the literature (see Quadrini and Ríos-Rull (2014), Cagetti and De Nardi (2008), and De Nardi and Fella (2017)), the model with a canonical earnings process is unable to generate the substantial level of wealth concen-

³⁰The earnings process used by Karahan and Ozkan (2013) is similar to our age-dependent process with normal shocks. Their estimate of 0.38 for the true coefficient ψ^p is in the ballpark of our estimate of 0.46 in Table 2.

	Percentage wealth in the top						
	Wealth						
	Gini	1%	5%	20%	40%	60%	80%
Data (SCF 1989)	.79	30	54	81	94	99	100
Model: Canonical	.64	9	29	65	88	97	100.1
Model: Nonlinear process	.61	7	25	61	85	96	99.9

Table 3: Wealth distribution

tration that we observe in the data. For instance, the top 1% of agents holds about 30% of total wealth in the data, while the corresponding share is only holds 9% in the model. Comparing the second and third rows in the table reveals that allowing for nonlinear earnings does not improve the fit of the wealth the distribution. If anything it marginally reduces the degree of wealth concentration at the top.³¹ One may be concerned that this may be due to the nature of the PSID data, which top-codes earnings and does not oversample the rich. However, [De Nardi et al. \(2016\)](#) conduct a similar exercise using synthetically-generated W2 Social Security Administration tax data, and find similar results for the concentration of wealth at the top. As pointed out by [De Nardi and Fella \(2017\)](#), non-homothetic preferences for bequests, entrepreneurship, and medical-expense risk are important for life-cycle models to be able to account for top wealth concentration.

2.7.4 Welfare

The differences in the evolution of the variance of log consumption and the pass-through of income shocks to consumption show that income risk affects households in a different way in the two economies. A natural question is to which extent these differences affect welfare.

To measure welfare, Table 4 displays the constant fraction of consumption that households are willing to give up to live in a world with no income uncertainty; i.e., a world where earnings are equal to the common and deterministic average earnings profile. We compute this measure under the veil of ignorance (before people enter the labor market and draw the first earnings realization) and, for comparability, we keep the discount factor the same for both processes fixing it at its calibrated value for the nonlinear process.

³¹We target a wealth to income ratio of 3.1, but this has little effect on wealth concentration.

	Welfare cost
Canonical process	28.3%
Nonlinear process	26.2%
Normal, age-dependent	24.3%
Non-normal, age-dependent	25.4%

Table 4: Consumption measure of welfare costs

The nonlinear process features larger variance and lower persistence of persistent shocks that, as we have discussed in Section 2.7.1, improve shock insurability, but also negative skewness and high kurtosis, as well as nonlinear persistence. Vice versa in the canonical model, the lower variance of shocks at all ages after the first one is counteracted by their high persistence (unit root) and the higher variance of the initial condition. The net effect of all these forces is that overall risk is higher under the canonical process. In particular, households would be willing to give up 28.3 per cent of their consumption in every state to eliminate earnings risk under the canonical earnings process compared to 26.2 per cent under the nonlinear earnings process.

In order to understand the respective contribution of the various features of our rich earnings process we have also computed the welfare cost under the two *intermediate* earnings processes considered in sections 2.7.1 and 2.7.2. The results are reported the third and fourth row of Table 4. They show that the lower welfare costs of earnings risk relative to the canonical process are all due to the age-dependence of second moments. Allowing only for age-dependence reduces the welfare costs of earnings risk by 4 percentage points, from 28.3 to 24.3 per cent, relative to the canonical process. Introducing, non-normality lowers welfare by one percentage point, relative to the normal age-dependent case, while allowing also for non-linearity reduces it by an additional percentage point.

2.8 Conclusions

We estimate a richer stochastic process for *household disposable earnings* featuring a transitory and persistent component and allowing for age-dependence, non-normality and nonlinearity. We use a standard life-cycle model with incomplete markets to compare the implications of our richer process to those of canonical permanent/transitory linear process with age-independent, normally-

distributed shocks. Our main findings are as following. Compared to the canonical process, the richer process implies a much better fit of the growth in cross-sectional consumption dispersion over the life cycle and a degree of self-insurance of persistent earnings shocks in line with the empirical estimates in [Blundell et al. \(2008\)](#). It also implies smaller welfare costs of earnings fluctuations. In terms of wealth inequality, we find that the two earnings processes have similar implications, including at the upper tail of the wealth distribution.

Based on this evidence, in Chapter 3 we now turn to describing how earnings dynamics have changed over different generations in the United States and its implications for homeownership and household portfolio decisions.

Chapter 3 Intergenerational Changes on Household Earnings Risk, Homeownership, and Household Portfolios

3.1 Introduction

The economic conditions faced by young US households are radically different from those that their parents and grandparents experienced when they were their age. Jobs are more unstable than they used to be, with career-long positions becoming less and less prevalent, and earnings inequality has increased. While the labor incomes of high earners have increased substantially over time in real terms, income-poorer individuals have seen their earnings stagnate or decrease.³²

Meanwhile, homeownership has shrunk. Within the cohort born in the 1940s, at age 35 almost 75% of households were living in houses they owned. The figure was ten percentage points lower for those born in the 1960s, and more than 20 percentage points lower for the early ‘Millennials’ born in the 1980s. This happens in a context in which financial markets have become more developed³³ and stock market participation has been increasing for younger generations.

This Chapter studies the role of these changes in household labor income dynamics and financial conditions in explaining homeownership and portfolio composition across generations. To do so, it proposes two novel contributions. First, it designs a flexible, cohort and business-cycle dependent earnings process, based on [Arellano et al. \(2017\)](#) and Chapter 2 of this thesis, that allows shocks to household labor income to be age-varying, non-normal, non-linear, and correlated with stock market returns and house prices, as in the data. Second, it builds and calibrates a rich life-cycle model with correlated aggregate and idiosyncratic risk, in which

³²These facts have been established in a large literature surveyed in [Acemoglu and Autor \(2011\)](#) and [Goldin and Katz \(2009\)](#). [Güvenen, Kaplan, Song and Weidner \(2017\)](#), using US administrative data, and [Borella, De Nardi and Yang \(2019\)](#), using survey data, find decreases in median male wages in real terms between the cohorts born in the 1940s and the cohorts born in the 1960s.

³³[Dynan, Elmendorf and Sichel \(2006\)](#) describe how financial deregulation, changes in risk-assessment methods, and the expansion of secondary markets increased the fraction of households with access to credit and how much those who already had access could borrow.

households decide their consumption, savings, housing stocks, portfolio share of safe and risky assets, and mortgage debt. Importantly, households only need to satisfy downpayment constraints and income tests at the time of mortgage origination, which implies that the outstanding mortgage can go above the value of the house if there is a negative shock to house prices. Households can also hold liquidity whilst they have a mortgage.

I use the model to compare the life experiences of three generations, namely, those born in the 1940s, 1960s, and 1980s. I assume that an American born in the 1940s differs from younger generations in three main ways. First, they face different experiences in the labor market. I use household earnings data from the Panel Study of Income Dynamics (PSID) to estimate the earnings process separately for all three generations, thus incorporating the changes in earnings inequality and earnings risk in a flexible, data-driven manner. I separate the persistent and transitory components of earnings, which allows me to control for potential measurement error in the survey. Second, they face different conditions in financial and housing markets. Housing has become more expensive over time with respect to average incomes, and different generations entered the labor market in different stages of the business cycle or the house price cycle. Third, the 1980s generation faced particularly looser financial constraints when they started to buy houses in the early 2000s, which I capture with a reduction in downpayment constraints.

Time, age, and cohort are explicit in the model. Average earnings, homeownership, and stock market participation at each age differ across generations as they do in the data. I do not homogenize age profiles across cohorts and thus do not need to disentangle year and cohort effects to obtain them.³⁴ I adopt the actual realizations of house prices and stock market returns each year from historical data, and use the Survey of Consumer Finances (SCF), including its earlier versions dating back to 1963, to obtain information about household portfolio compositions by age and generation.

³⁴Age, year, and cohort are collinear. To obtain age profiles in a sample with several cohorts and years, the usual practice is to either remove year fixed effects or cohort fixed effects, which can lead to very different implications. See [Heathcote, Storesletten and Violante \(2005\)](#) for a discussion on how the choice of removing year or cohort effects impacts measures of earnings and consumption inequality, and [Ameriks and Zeldes \(2004\)](#) for the effect on household portfolio shares.

The main results are as follows. First, intergenerational changes in earnings dynamics, asset returns, and housing prices obtained from the data fully explain the differences in homeownership between the 1940s and 1960s cohorts. For the 1980s cohort, who started to buy houses in the early 2000s, looser borrowing constraints partially counteracted the effect of high house prices. I do not need to assume that preferences have changed to explain the lower homeownership rates for younger generations.

To isolate the effect of changes in labor market income dynamics, I perform a counterfactual experiment in which I attribute the earnings process of the 1940s cohort to the younger generations, whilst keeping all other elements of the model constant, including house prices. More than half of the difference in homeownership at age 30 for both generations can be accounted for by changes in earnings inequality and risk. Not all of it is due to delayed home-buying: changes in earnings dynamics still have an important effect at age 40 and afterwards. These results are robust to letting house prices adjust, assuming an empirically plausible level of housing supply elasticities.

The main driver of these changes is the increase in earnings inequality at labor market entry, with a more limited role for the increase in earnings risk. The intuition is simple. Households with lower initial and expected lifetime earnings find it harder or suboptimal to engage in a large expenditure like a house, which would leave them with a sizeable mortgage with respect to their current income, and thus exposed to income and house price risk.

Second, the increase in stock market participation of younger cohorts can be rationalized with a substantial reduction in stock market participation costs, which reflects easier information acquisition. Today, many workers who are starting new jobs either receive information about retirement accounts or are automatically enrolled into retirement plans like IRAs or 401(k)s.

These intergenerational changes also have implications for household wealth accumulation. Although financial wealth represents an increasing share of household portfolios for those born after the 1970s, it is more unequally distributed than housing wealth. In the 1940s genera-

tion, relatively poorer households who wanted to be homeowners bought a house, got a mortgage, and benefitted from gains in the housing market. Similarly ranked households in younger generations are no longer buying houses, and they do not fully compensate the lack of housing wealth by saving in financial assets. The model predicts that lowering the cost of access to financial markets for lower and middle income households can help to increase their wealth holdings and thus reduce wealth inequality. Additionally, these changes in earnings dynamics and household portfolio composition impact the way households react to aggregate and idiosyncratic shocks. The model suggests that younger generations display larger consumption responses to persistent income shocks, which can have important implications for monetary and fiscal policy design.

Overall, these findings suggest that changes in labor market income dynamics and in the housing market are having substantial effects in the life experiences of most Americans, and they can influence, in the longer term, the distribution of income and wealth, intergenerational mobility, and the effects of policies.

Related literature

In this Chapter I introduce a flexible process for earnings dynamics, which I input into a rich model of housing and portfolio choice over the life cycle to understand intergenerational changes and their macroeconomic implications. Thus, I contribute to three main strands of the literature.

Earnings dynamics. The earnings process that I propose in this paper jointly considers rich features of earnings risk, business cycle variation, and changes over the generations. As such, it expands on a broad literature on earnings dynamics, and its effects on consumption, welfare, and portfolio allocations³⁵.

A number of recent contributions have documented that earnings risk varies by age and by the position of an individual in the earnings distribution, and that earnings shocks are left-skewed and leptokurtic (e.g. [Guvenen et al. \(2016\)](#) with US administrative data for individuals). [Arel-lano et al. \(2017\)](#) devise an econometric framework that allows for the

³⁵See [Meghir and Pistaferri \(2011\)](#) for a summary of this branch of the literature.

separate identification of the distributions of the persistent and transitory components of earnings whilst allowing for flexibility in their distributions, and thus accommodating all of these non-normal and non-linear features. Using their framework, in the second chapter of this thesis I introduce a flexible earnings process, with a persistent and a transitory component, into a standard life cycle model. I find that allowing for these rich earnings dynamics helps to better understand the evolution of cross-sectional consumption dispersion and the extent to which households can self-insure against persistent earnings shocks.

However, these estimated processes abstract from business cycle variation. [Storesletten et al. \(2004a\)](#) show that, in the context of a standard earnings process with normal shocks, the standard deviation of earnings fluctuations is strongly countercyclical. [Guvenen, Ozkan and Song \(2014\)](#) argue that the key element that fluctuates over the cycle is the left-skewness of earnings shocks: during recessions, large drops in earnings become more likely. This business cycle component of earnings risk and its correlation with asset returns is important to understand household portfolio decisions. In this Chapter, I propose an extension of the econometric framework devised by [Arellano et al. \(2017\)](#) that allows for business cycle variation in earnings dynamics in the form of a Markov-switching regime, and that displays, when estimated in survey data, the rich features described in [Guvenen et al. \(2014\)](#).

Another recent contribution that designs and implements an earnings process with variation in higher order moments over the business cycle is [Busch and Ludwig \(2017\)](#). Both their approach and their focus differ from mine. I use a flexible nonparametric model that I estimate in panel data, while they define a rich parametric process and estimate it, à la [Storesletten et al. \(2004a\)](#), by using cross-sectional moments identify the sequence of past shocks. I focus on the relationship of rich earnings risk with changes in household portfolio compositions, while they study the welfare costs of risks. Furthermore, my approach also allows for variations in earnings dynamics over different cohorts.³⁶

³⁶[Lippi and Perri \(2019\)](#) show that the changes in household income dynamics in the US over the past 50 years have an important role in explaining the evolution of inequality and part of the reduction in aggregate growth.

Housing and portfolio choice over the life cycle. An extensive literature has studied the determinants of housing demand over the life cycle, its relationship to nondurable consumption and savings, and its interaction with household responses to income or house price shocks (Atanasio, Bottazzi, Low, Nesheim and Wakefield (2012), Berger, Guerrieri, Lorenzoni and Vavra (2017), etc.). Houses are a large part of the portfolio of most households, and passive saving through house price appreciation is an important determinant of wealth accumulation (Fagereng, Holm, Moll and Natvik, 2019). Additionally, houses have a preferential tax treatment in most countries, which occurs both because owner-occupied rents of housing are not taxed and because of government programmes like the US mortgage interest tax deductibility (Gervais (2002), Díaz and Luengo-Prado (2008), Nakajima (2010), etc.). This study incorporates all of these important dimensions in the modelling of houses.

In parallel, many papers have used models of portfolio choice over the life cycle to explain important puzzles, such as the high equity premium or the low level of stock market participation. Standard asset allocation models over-predict how many people invest in stocks and, conditional on participation, how much of their wealth they invest in them. This puzzle can be overcome considering alternative preferences and costs of participation in the stock market (Gomes and Michaelides (2005), Alan (2006)), or the correlation between labor market income risk and stock market risk, although its effect is usually quantitatively small.³⁷ I incorporate these preferences, costs, and correlations to my modeling of household portfolio decisions.

Fewer contributions have explored, like I do, the interaction between housing, portfolio choice, and the life-cycle.³⁸ Cocco (2005) shows that younger and poorer investors have less financial wealth to invest in stocks because they prefer to start investing in housing, and that this reduces the benefits of equity holdings, thus helping solve the stock market partic-

³⁷See Ameriks and Zeldes (2004), Cocco, Gomes and Maenhout (2005), Benzoni, Collin-Dufresne and Goldstein (2007) and Fagereng, Gottlieb and Guiso (2017). Fagereng, Guiso and Pistaferri (2016) argue that this small quantitative effect arises from problems in the identification of uninsurable risk.

³⁸These include Flavin and Yamashita (2011), Yao and Zhang (2005), and Vestman (2012), who focuses on the role of preference heterogeneity to explain why homeowners participate more in the stock market. Becker and Shabani (2010) and Chetty, Sándor and Szeidl (2017) study the role of mortgage debt on portfolio allocations.

ipation puzzle. My life-cycle model is similar to his, in that it allows for housing and portfolio choice decisions, but we differ in our focus. While his paper studies how housing crowds out stock market participation, I focus on the joint role of housing and portfolio choice in life-cycle wealth accumulation, the role of labor market income risk, and intergenerational changes. Furthermore, my model is richer and includes flexible earnings risk, mortgages that do not need to satisfy LTV constraints in every period, the possibility of renting, and a richer process for stock returns that features a disaster state.

Intergenerational changes: cohort and time effects. Changes over time and over the business cycle in asset returns, house prices, and labor market dynamics affect both the decision to buy a house and the allocation between safe and risky assets. The link between those and macroeconomic outcomes is still relatively unexplored.³⁹ Closely connected with this study are [Fisher and Gervais \(2011\)](#), who in a stationary equilibrium framework find that the increase in earnings uncertainty is a major candidate to explain the reduction in homeownership of the young between 1980 and 2000. This study builds on their contribution along several dimensions. First, I explicitly consider intergenerational differences by modelling each cohort separately, which allows me to better capture cohort and year effects on earnings and asset prices, including variations in price to income ratios of housing. Second, in my model house prices are risky and agents can hold liquidity while they have a mortgage. Both are important elements because they affect the risk associated with buying a house: the former increases household exposure to risk, but the latter decreases it, because it allows them to better smooth income fluctuations. Third, I study the role of housing in the context of a richer household portfolio decision, and thus can accommodate possible substitution effects across asset classes as housing prices and asset returns change over time.

³⁹Relevant contributions include [Nakajima \(2005\)](#), who suggests that rising earnings inequality in the U.S. can be related with the increase in housing prices and lower return of financial assets, or [Chambers, Garriga and Schlagenhaut \(2009\)](#), who study the boom in homeownership between 1994 and 2005 and relate it to mortgage innovations. [Fischer and Khorunzhina \(2019\)](#) relate changes in homeownership rates over the life-cycle to increases in divorce rates, that trigger precautionary savings for the young but reduce homeownership for older households.

3.2 An overview of intergenerational changes

The 1940s, 1960s, and 1980s generations have experienced different economic environments, both in terms of their labor market experiences and the returns to their assets, and taken different economic decisions with respect to buying houses and investing in stocks. I now turn to empirical evidence to describe these differences in detail.

3.2.1 Distribution of earnings

The income of the median earner at each age differs across generations. For men, Social Security data shows that median income at labor market entry increased between the generations born in the 1937 and 1947, but has decreased for those thereafter (Güvenen et al., 2017). As the shape of the life-cycle profile of earnings has changed little, both median earnings at each age and median lifetime earnings are lower for younger generations.

The top left panel of Figure 22 shows the corresponding profiles for the PSID, deflated using the CPI.⁴⁰ Consistently with the administrative data, the earnings of the median male earner in the PSID at age 25 have decreased from the cohort born in the 1940s, which entered the labor market in the early 1960s, to the cohort born in the 1960s by around 12% in real dollars. However, during this period there was a significant increase in female labor force participation and women’s wages, which acted as a counteracting force and almost completely reversed this decrease in terms of household earnings (bottom left panel). After age 30, when most household formation has taken place, median household earnings are higher for the younger cohort than for those born in the 1940s. Deflating using the PCE generates even larger differences across cohorts (see Appendix 3.3.2).

The decrease in earnings is also less clear if we look at averages rather than medians (central top and bottom panels), which suggests that earnings have become progressively more right-skewed and the earners above the median have seen larger increases than the earners below the median.

The two right-hand side panels of Figure 22 confirm this intuition and show that earnings dispersion has grown for younger cohorts, particularly

⁴⁰Appendix C describes the data and sample selection procedures I use, and Appendix 3.3 shows robustness with respect to these choices.

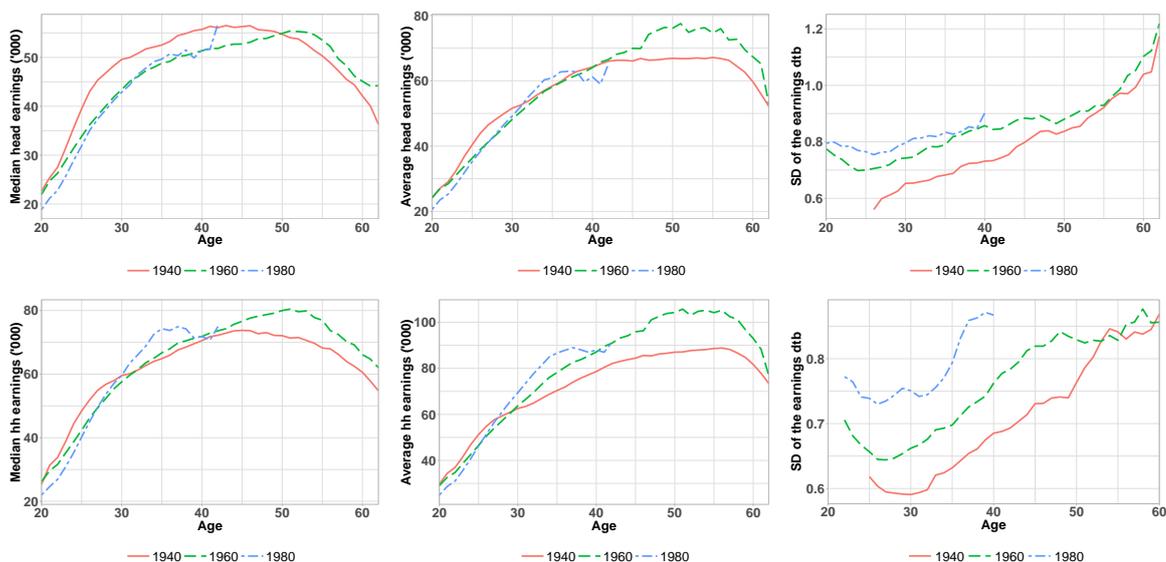


Figure 22: Changes in the earnings distribution over the generations. Top: household heads; bottom: household earnings. Left: median earnings, center: average earnings, right: standard deviation of the log earnings distribution. PSID data.

in terms of household earnings. Most of the difference is already present at age 25. This large increase, together with little action in the means, implies that the earnings-poorest of more recent cohorts are relatively worse off than people in the same percentile of earnings of earlier cohorts, and conversely the earnings-richest are better off today. This has important implications for the timing and features of their homeownership decisions, as I describe when I turn to my model and its results in Section 3.6.

Part of these changes (in particular, the reduction in median and average earnings at younger ages) can reflect intergenerational changes in family composition, and, in particular, delayed household formation. Appendix 3.3.3 shows that, if we restrict the sample to married couples, all patterns are consistent with the main picture, in particular in terms of earnings inequality. I consider this as suggestive evidence that the timing of family formation is not the only driver of the transformations we observe.

3.2.2 Earnings risk

Apart from a more spread earnings distribution, younger cohorts also face increasing earnings risk. Figure 23 shows that, in general, the standard deviation of earnings *changes* is larger for younger cohorts. The most sig-

nificant differences between the 1940s and 1960s cohort are concentrated between ages 30 and 50, and the 1980s cohort started its working life with a very large level of earnings variability. However, those changes do not extend to higher-order moments of earnings risk. As Figure 24 shows, earnings changes display negative skewness and high kurtosis for all cohorts, but there has been little change in those measures over time⁴¹. However, this observation does not imply that the tails of all three distributions are equally fat: given a level of kurtosis and skewness, increasing the variance makes large shocks more likely than before.

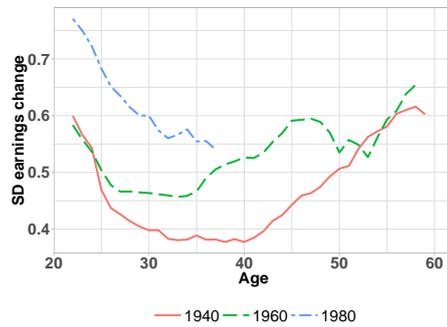


Figure 23: Standard deviation of log earnings changes, by cohort

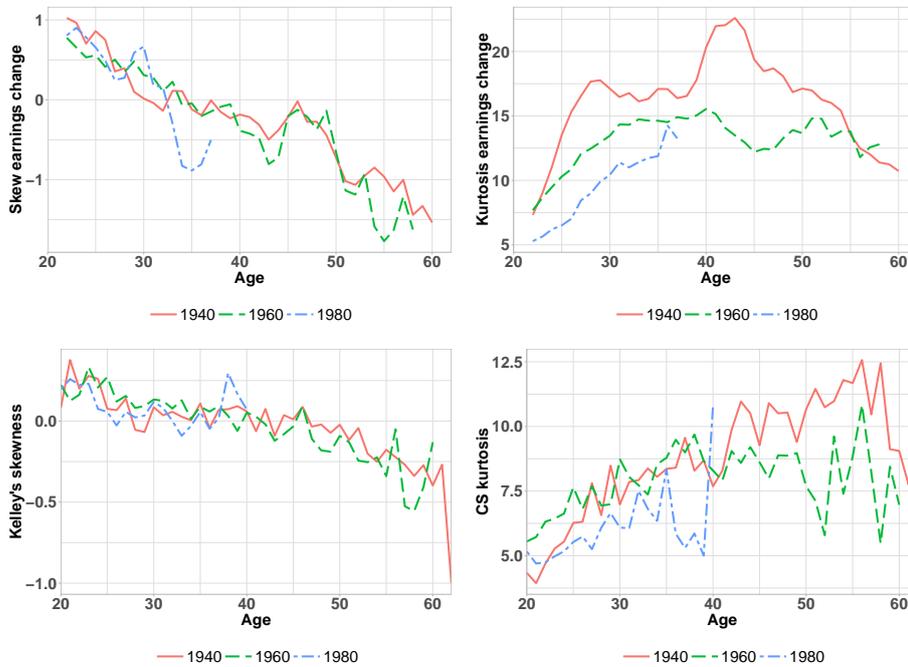


Figure 24: Higher order moments of log earnings changes, by cohort

⁴¹Appendix 3.1.3 provides definitions for these.

3.2.3 Housing prices and stock returns

The ratio of median house prices to median income has increased, on average, in the United States over the last 60 years. Younger generations, at the same age, now have to devote more years of their income to buy a home compared with their parents.

The left panel of Figure 25 shows the evolution of median price-to-income (PTI) ratios, based on PSID data⁴², from 1975 to 2017. Two main features, which are reproduced in the model, are particularly salient. First, PTI ratios have been increasing over time. Second, there are large cyclical variations in house prices, although they are not always correlated with the business cycle. These induce an additional source of variation across cohorts, as some of them may have entered the labor market in a time where house prices were cyclically low, and benefitted from the situation to make housing purchases earlier on in their lives.

On the other hand, the evolution of stock returns (right panel of Figure 25) shows large fluctuations, which are more strongly correlated with the business cycle, but fewer secular trends.

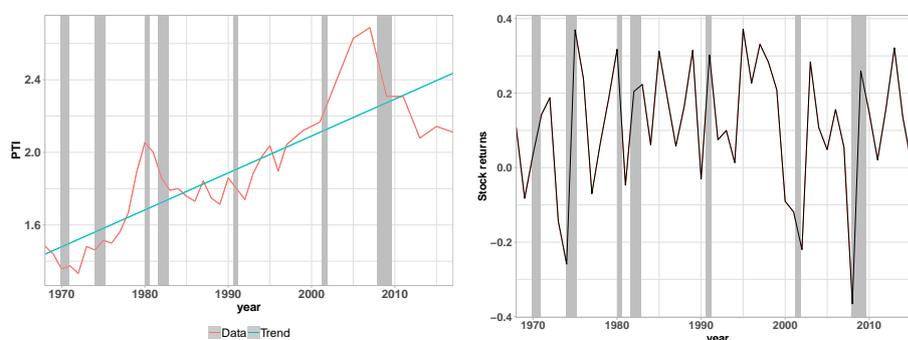


Figure 25: Evolution of asset prices and returns. Left: Price to Income ratios for housing (PSID data). Right: Stock returns, S&P 500. Shaded areas correspond to NBER recessions.

3.2.4 Financial conditions

The process of financial deregulation and innovation that started in the 1980s and expanded during the 1990s improved the access of households

⁴²Lovenheim (2011) shows that both median and mean home price indices constructed from PSID data track Federal Housing Finance Agency repeat home sales indices very well.

to credit, both from an extensive (more people can get credit) and intensive (the same household can borrow larger amounts) perspective. See, for instance, [Gerardi, Rosen and Willen \(2007\)](#) for a detailed description of the regulatory changes, the changes in the structure of the financial sector, and the new mortgage products that became available over this period. These changes were partially encouraged by policymakers, who were worried about low homeownership rates (a salient example is Bill Clinton’s National Homeownership Strategy).

Another important change was the introduction of tax-advantaged retirement accounts, such as individual retirement accounts (IRAs), which started in 1974 and became popular in the 1980s, and 401(k)s, which were introduced in 1978 and also became popular later on. Later reforms made these accounts more beneficial and less restricted, and automatic enrollment in pension plans further increased the number of stock market participants by reducing both the financial and psychological costs of enrollment.

3.2.5 Homeownership and portfolio composition

Parallel with the changes described earlier, homeownership rates have been falling for recent cohorts. I use the word homeownership to refer to the percentage of households that live in owner-occupied housing - this differs from its alternative, more common usage of the percentage of homes that are occupied by their owners.

Using PSID data (Figure 26, top panel), we observe that, at age 35, homeownership has dropped by over 10 percentage points between the cohorts born in 1940 and 1960, and by another 10 percentage points between the cohorts born in 1960 and 1980⁴³.

At the same time, stock market participation has increased significantly for younger cohorts (Figure 26, bottom left). This is related to the introduction and generalization of retirement accounts I have just described in Section 3.2.4, an explanation which is reinforced by the small differences across cohorts in direct stock market participation (see Figure C.9

⁴³The picture in terms of intergenerational differences is similar under alternative sample selection procedures (Appendix 3.3.1) and considering only married households or households with children (Appendix 3.3.3).

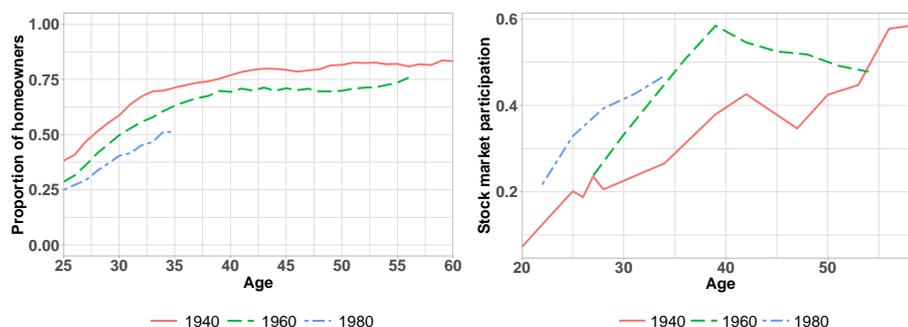


Figure 26: Homeownership (left) and stock market participation (right), three generations. PSID and SCF data.

in Appendix 3.3.6).

However, stock market participation also seems to display strong year effects. For instance, direct stock market participation increased significantly in the years before the 2000 stock market crash, and dropped dramatically afterwards, as it can be seen in the profile for the 1960s cohort when they were 40 years old.

3.3 A business-cycle dependent earnings process

In this section I develop a flexible earnings process that can capture the differences across generations I have just described, whilst encompassing a set of elements that have been shown to be important to describe the features of household earnings risk and its implications on household consumption and self-insurance (De Nardi et al., 2019). These include age-varying persistence, variance, and higher order moments, non-normalities such as high negative skewness and large kurtosis, and non-linearities such as previous-earnings-dependent persistence.

The process is based on Arellano et al. (2017), but, on top of that, it includes three important factors: business cycle variation in earnings dynamics, including its non-normal and nonlinear features, intergenerational changes in the distribution of earnings, and intergenerational changes in earnings risk. The former is necessary because idiosyncratic risk correlates with aggregate asset price risk, which can have implications for household portfolio decisions and insurance over the business cycle. The latter two are necessary to address the questions posed in this study.

Let \tilde{y}_{it} denote the logarithm of pre-tax labor earnings, net of age effects, for household i of cohort c_i ($c_i \in \{1940, 1960, 1980\}$) living in calendar year t with age age_{it} . I assume earnings are the sum of a persistent and a transitory component:

$$\tilde{y}_{it} = \eta_{it} + \epsilon_{it} \quad (19)$$

where both have absolutely continuous distributions. The persistent component η_{it} is assumed to follow a first-order Markov process, while the transitory component ϵ_{it} has zero mean and is independent over time and of the persistent component.

We can introduce these assumptions by writing the processes for η and ϵ , and the initial condition for the persistent component η_1 as:

$$\eta_{it} = Q_\eta(\nu_{it}^\eta | \eta_{i,t-1}, age_{it}, c_i, \Omega_t^y), \nu_{it}^\eta \stackrel{iid}{\sim} U(0, 1), t > 1 \quad (20)$$

$$\epsilon_{it} = Q_\epsilon(\nu_{it}^\epsilon | age_{it}, c_i), \nu_{it}^\epsilon \stackrel{iid}{\sim} U(0, 1) \quad (21)$$

$$\eta_{i1} = Q_{\eta_1}(\nu_{i1}^{\eta_1} | age_{i1}, c_i, \Omega_1^y), \nu_{i1}^{\eta_1} \stackrel{iid}{\sim} U(0, 1) \quad (22)$$

Equation 20 specifies the dependence of η_{it} on its previous realization with a flexible quantile function Q_η . This function depends on the age of the household, age_{it} , its cohort, c_i , and the aggregate state of the labor market, Ω^y . Thus, the features of earnings shocks are allowed to be different in expansions and recessions.⁴⁴ In this way, this formulation explicitly includes age, cohort, and year effects.

Q maps draws ν_{it} from the uniform distribution $U(0, 1)$ into quantile draws for η . ν_{it} can be thought of as a rank: if it is 0.9, it implies that the realization of η_{it} is on the 90th percentile conditional on age and $\eta_{i,t-1}$. A similar reasoning follows for the initial realization of the persistent component, with the further simplification that it only depends on age, cohort, and the current state of the labor market; and for the transitory component, which only depends on age and cohort. I treat the transitory component as measurement error or alternatively as a fully-insurable source

⁴⁴Section 3.4.3.1 describes the aggregate state in more detail, and Section 3.5.1 explains its implementation from the data.

of earnings fluctuations.

Following [Arellano et al. \(2017\)](#), to estimate the process I specify a parametric form for the quantile functions as low order Hermite polynomials:

$$Q_\eta(q|\eta_{i,t-1}, age_{it}, c_i, \Omega_t^y) = \sum_{j=0}^J a_j^\eta(q, c_i, \Omega_t^y) \psi_j(\eta_{i,t-1}, age_{it}) \quad (23)$$

$$Q_{\eta_1}(q|age_{i1}, c_i, \Omega_1^y) = \sum_{j=0}^J a_j^{\eta_1}(q, c_i, \Omega_1^y) \psi_j(age_{i1}) \quad (24)$$

$$Q_\epsilon(q|age_{it}, c_i) = \sum_{j=0}^J a_j^\epsilon(q, c_i) \psi_j(age_{it}) \quad (25)$$

where the coefficients a_j^i , $i = \epsilon, \eta_1, \eta$, for all states are modelled as piecewise-linear splines on a grid $\{q_1 < \dots < q_L\} \in (0, 1)$.⁴⁵ The intercept coefficients $a_0^i(q)$ for q in $(0, q_1]$ and $[q_L, 1)$ are modelled as the quantiles of an exponential distribution with parameters λ_1^i and λ_L^i respectively. All coefficients are allowed to differ across cohorts.

If one could directly observe the two components ϵ_{it} and η_{it} , it would be possible to find the coefficients above by quantile regression at each point of the quantile grid q_j . However, both components are latent. To deal with this, the estimation starts at an initial guess for the coefficients and iterates between draws of the posterior distribution of the latent persistent components and proceeds to find the coefficients by quantile regression. The process is repeated until convergence of the sequence of coefficient estimates.

This process nests more standard earnings process such as that proposed in [Storesletten et al. \(2004a\)](#), which I refer to as *canonical process*:

$$y_{it} = \eta_{it} + \epsilon_{it} \quad (26)$$

$$\eta_{it} = \rho\eta_{it-1} + \xi_{it} \quad (27)$$

with $\xi_{it} \sim N(0, \sigma_\xi^2)$, $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$ and

⁴⁵Following [Arellano et al. \(2017\)](#), I use tensor products of Hermite polynomials of degrees (3,2) in $\eta_{i,t-1}$, and age for each state k of $Q_{\eta,\Omega}(q|\eta_{i,t-1}, age_{it})$ and second-order polynomials in age for $Q_\epsilon(q|age_{it})$ and $Q_{\eta_1,\Omega}(q|age_{i1})$.

$$\sigma_t^2 = \begin{cases} \sigma_{r,c}^2 & \text{if } \Omega^y = \text{Recession} \\ \sigma_{b,c}^2 & \text{if } \Omega^y = \text{Boom} \end{cases} \quad (28)$$

where usually $\sigma_{r,c}^2 > \sigma_{b,c}^2$. Unlike in this process, my procedure implies that there is no need to assume age-independence or normality of earnings shocks, nor linearity in the dependence of the persistent component on its past realizations. While the earnings process is estimated on pre-tax rather than post-tax household earnings, most of its features regarding non-linearity and non-normality are qualitatively similar to [De Nardi et al. \(2019\)](#) and therefore I refer the interested reader to the discussion therein. Furthermore, the earnings process I propose can accommodate business-cycle varying features of higher order moments of earnings risk, such as countercyclical skewness.

I estimate the earnings process on PSID data for all three cohorts. Given that the PSID became biennial from 1997 onwards, the period is two years for both the earnings process and the structural model. I use the full length of the PSID (1968-2017).⁴⁶ More details about the data treatment, cohort definitions, and sample selection are available in [Appendix C](#).

3.3.1 Implications of the earnings process

3.3.1.1 Intergenerational differences

The earnings process captures the intergenerational changes in earnings dynamics documented in [Section 3.2.2](#) well (see [Appendix 3.3.4](#)). An additional notion that has changed over time is *nonlinear persistence* ([Figure 27](#)) by previous earnings and the quantile of the earnings shock. For the youngest cohort, persistence is much larger for higher-income agents and all ranks of their shocks, and lower for low-income agents, particularly for large shocks.

⁴⁶The semiparametric implementation of the nonparametric model defined in [Arellano et al. \(2017\)](#) allows to interpolate and obtain an earnings process for every state and age even if not all combinations are present in the data.

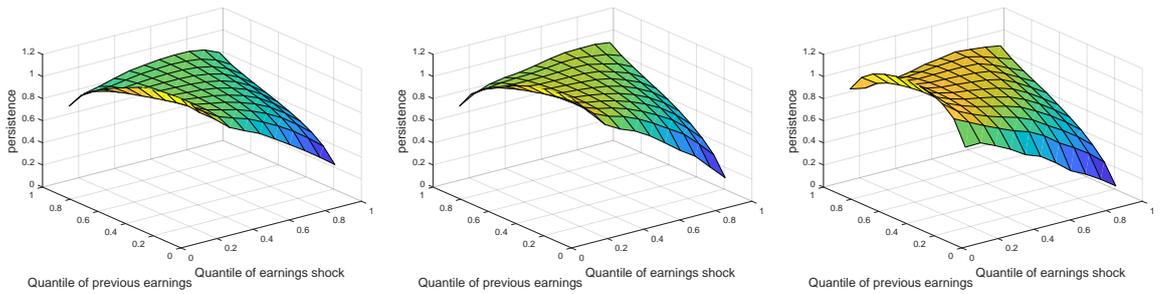


Figure 27: Nonlinear persistence, by cohort. Top left: 1940s; top right: 1960s; bottom: 1980s. For all cohorts: all agents below 40 years old. Figure C.8 in Appendix 3.3.5 provides the version with all ages included but thus varying composition across cohorts.

3.3.1.2 State-dependence of the earnings process

The state dependence of the earnings process implies that it has potentially different features in expansions and recessions. The left panel of Figure 28 shows the average expected change in earnings for individuals in different points of the earnings distribution for both aggregate states for the 1940s cohort. The other three cohorts display similar qualitative characteristics. During normal times, most individuals expect slight increases in their earnings. The very poorest expect the highest improvements in relative terms, while there is a certain level of mean reversion for the earnings-richest. In recessions, the expected increase in earnings shifts downwards. The earnings-poorest expect lower increases, and the earnings-richest expect larger drops.

These average measures mask significant heterogeneity. Figure 28, right panel, plots Kelley’s measure of skewness of earnings changes during an expansion and during a recession. During normal times, the skewness is basically zero for most of the distribution: the distribution of earnings changes is symmetric and large negative shocks and large positive shocks of equal magnitudes are equally likely. However, during a recession skewness becomes negative, particularly so for the very richest. This implies that large decreases in earnings become more likely with respect to large increases in earnings.

Capturing these features of the distributions is important to better understand how households take portfolio decisions. For instance, the combination of high likelihood of disaster risk in the stock market and

large skewness in labor earnings for a particular household can explain why they choose to keep some of their savings in safer investments.

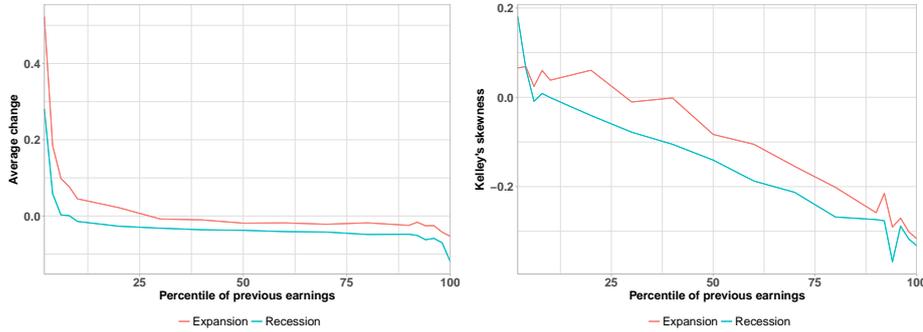


Figure 28: Average expected change in earnings (left), and Kelley’s skewness of earnings changes (right), by previous earnings percentile and aggregate state of the economy. Percentiles refer to the distribution during an expansion.

An additional realistic feature that the Markov-switching earnings process captures is history dependence: at any point in time, the distribution of earnings for a given cohort depends on the set of expansions and recessions that the cohort has lived through. In particular, the recovery from recessions is usually sluggish.⁴⁷ Figure 29, left panel, shows that the earnings process I propose replicates this feature without large increases in the state space. It represents, for the simulated earnings process of the 1940s cohort, the percentage difference in average earnings between a cohort that underwent a single recession at age 44 (“NL process”) and one that never lived through a recession throughout its entire labor market history. Suffering one recession has important effects on impact that last for relatively long. In contrast, the canonical earnings process generates a counterfactual *increase* in average earnings because higher variances in logs, at a constant average, imply higher averages in levels. The purple line represents it for the estimated parameters in Storesletten et al. (2004a), while the yellow line represents the canonical counterpart to the process I estimate for the 1940s cohort (see Appendix 3.4.3 for details).

The rich earnings process also captures differential impacts by initial position in the earnings distribution (middle panel). Recessions affect the

⁴⁷A broad literature has studied both the large negative long-run effects of displacement for individual workers (e.g. Jacobson, LaLonde and Sullivan (1993)), that are particularly severe within recessions (Davis and Von Wachter, 2011), and the slow recovery of employment after downturns like the Great Recession (Ravn and Sterk, 2017).

earnings of the highest and lowest earners by more than those around the median. The impact of a recession also differs by age (right panel): younger agents are hit harder and take longer to recover. By construction, the canonical earnings process does not replicate either of these facts.

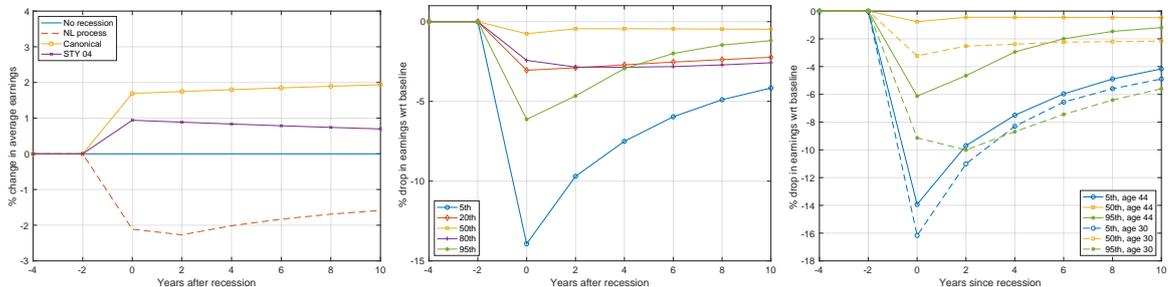


Figure 29: Earnings by age with respect to the counterfactual in which a recession never occurs. Left: average earnings, recession at age 44. Middle: by initial earnings percentile, recession at age 44. Right: by initial earnings percentile, comparing ages 30 and 44.

3.4 Model

I build a life-cycle structural model to evaluate to which extent the changes in earnings and financial conditions described in Section 3.2, modelling the former using the process described in Section 3.3, can account for the intergenerational differences in homeownership and portfolios I described earlier.

In the model, the economy is populated by a continuum of households i that belong to cohort c . From the perspective of a cohort, age and time are equivalent and indexed by t . The model period is two years. Households enjoy nondurable consumption and housing, are subject to exogenous earnings risk, and can hold three types of assets:

- Safe, liquid accounts.
- Housing (if they don't hold any, they must participate in the rental market).
- Risky financial assets, which they cannot short.

To finance their housing expenditures, they can also hold liabilities in the form of mortgages.

3.4.1 Demographics

Households are born in the model at age 20, retire at age 60 and face positive and increasing death probabilities ξ_t starting at that age. They die for sure at age 86. An average demographic profile at each age is introduced in the model with a taste shifter θ_t , which represents the average OECD equivalence scale at each age, and generates age-varying marginal utility from nondurable and housing consumption.

3.4.2 Preferences

Preferences are [Epstein and Zin \(1989\)](#) and allow to disentangle the elasticity of intertemporal substitution ψ and the risk aversion coefficient γ . Whenever $\gamma > \frac{1}{\psi}$ (which is the benchmark case in this Chapter), they imply that agents prefer an early resolution of uncertainty, as standard in studies on the equity premium and risk-free rate puzzles, and in portfolio choice models ([Cocco et al. \(2005\)](#), [Campanale, Fugazza and Gomes \(2015\)](#) or [Kaplan and Violante \(2014\)](#)).

Utility at age t is therefore represented by:

$$U_{it} = [(\theta_t c_{it}^\nu s_{it}^{1-\nu})^{\frac{\psi-1}{\psi}} + \beta (\mathbb{E}_t U_{it+1}^{1-\gamma})^{\frac{1}{1-\gamma} \frac{\psi-1}{\psi}}]^{\frac{\psi}{\psi-1}} \quad (29)$$

where θ is the taste shifter described earlier, c is nondurable consumption, and s is the housing service flow. In this specification, β is the discount factor, ψ is the elasticity of intertemporal substitution, γ is the coefficient of relative risk aversion, and ν measures the relative importance of nondurable consumption with respect to housing. This Cobb-Douglas specification assumes an elasticity of substitution between housing and nonhousing of 1, which is justified by the almost constant shares of expenditure in housing in micro data (e.g. [Davis and Ortalo-Magné, 2011](#)). In practice, since housing in the model is discrete, this is equivalent to assuming that housing utility is a proportional scaling of the utility from nondurable consumption.

The utility value of housing s_t depends on the quality of the owned home and does not vary with its price. It is highest for owners of high-quality houses (\bar{s}^2), lower for owners of low-quality houses (\bar{s}_1), and lowest

for renters (\bar{s}^0).

Households value bequests left according to:

$$v(b) = \phi_1(\phi_2 + b)^{\frac{(\psi-1)}{\psi}} \quad (30)$$

This specification mimics, in an Epstein-Zin framework, [De Nardi \(2004\)](#). The term ϕ_1 determines the intensity of the bequest motive and ϕ_2 determines the extent to which bequests are a luxury good.

3.4.3 Environment and technologies

3.4.3.1 Aggregate state

During each year t , the economy is in an aggregate state Ω_t composed of three elements: the state of the housing market Ω^h , which determines house prices, the state of the stock market Ω^f , on which stock returns depend, and the state of the labor market Ω^y , which determines the evolution of the earnings process. Thus, $\Omega_t = \{\Omega_t^f, \Omega_t^h, \Omega_t^y\}$. Households know the process governing the aggregate state, and use it to make predictions about the future, which in turn affect their decisions.

Ω^h is a Markov chain of order 2. In the data, not only house prices are persistent, but also their *growth* is. The Markov 2 assumption allows agents to be aware of whether house prices are in an increasing or decreasing regime, which together with their current realization helps them to predict how they will evolve in the future.

The state of the labor market Ω^y is a Markov chain of order 1. Households observe whether the economy is currently in an expansion or in a recession, which helps them to predict the performance of the labor market in the following period. On the other hand, stock returns Ω^f are independent across periods.

However, the model allows the realizations of each of the three elements to be interdependent according to their empirical correlations. For instance, in a recession it is more likely that stock returns are lower, which implies that there is higher probability of a bad realization of Ω_{t+1}^f if there is a bad realization of Ω_{t+1}^y too. The assumptions regarding the structure of this correlation, and the estimation of the aggregate state in the data,

are described in more detail in Section 3.5.1.

3.4.3.2 Earnings

Log earnings are composed of a deterministic component, which depends on age, and a stochastic persistent component η_{it} , which depends on the aggregate state of the labor market:

$$\log y_{it} = f(t) + \eta_{it}(\Omega_t^y) \quad (31)$$

Section 3.3 contains more details about the earnings process and its estimation. Transitory shocks may be reflecting measurement error or almost fully insurable fluctuations, so to save on computational costs I do not include them in the model.

3.4.3.3 Liquid accounts

Liquid accounts a_t are risk-free and they yield an exogenous and constant interest rate r^a . They cannot be negative: if they wish to borrow, households must apply for a specific type of financial asset, mortgages m_t , which I describe in detail in Section 3.4.3.6.

$$a_{t+1} \geq 0 \quad (32)$$

3.4.3.4 Risky financial assets

Households can also hold risky financial assets or stocks f . Stock returns r_t^f depend on the aggregate state of the stock market Ω_t^f . Households cannot short financial assets, thus the constraint for stocks is:

$$f_{t+1} \geq 0. \quad (33)$$

When $f_{i,t} = 0$, households pay a fixed entry cost κ^f to start investing in stocks.⁴⁸ This cost represents psychological, financial, and technical barriers to start investing in the stock market (opening financial accounts, acquiring information about them, etc.), and is frequently used in the port-

⁴⁸Section 3.5.5 shows that results are robust to assuming per-period participation costs instead.

folio choice literature (Gomes and Michaelides, 2005). Once a household participates, there are no additional costs of adjusting financial assets.

3.4.3.5 Housing

Households can buy houses h that come in discrete sizes:

$$h_{i,t} = \{0, h_1, h_2, \dots, h_H\}, \quad (34)$$

where 0 indicates renting and the other values indicate increasing qualities of housing. The discrete specification for housing follows Attanasio et al. (2012). I set $H = 2$ due to computational considerations.⁴⁹

Average house prices p_t^h depend on the aggregate state of the housing market Ω_t^h . They are expected to grow, but fluctuate around a trend as described in Section 3.5.1. The price of the different housing qualities h_j is assumed to be a fixed fraction of average house prices $p_t^h(\Omega_t^h)$ ⁵⁰, which I denote h^j :

$$h^j = \frac{p_t^{h_1}(\Omega_t^h)}{p_t^h(\Omega_t^h)} \quad (35)$$

In practice, this implies that the price of low-quality houses is a fixed fraction of the price of high-quality houses:

$$p_t^{h_1}(\Omega_t^h) = \frac{h^1}{h^2} p_t^{h_2}(\Omega_t^h), \quad (36)$$

Housing is illiquid. Households pay a proportional transaction cost to buy or sell housing $\kappa^h p_t^{h_j}(\Omega_t^h)$, which depends on the price of the house which is being bought (as in Bajari, Chan, Krueger and Miller (2013)). It reflects the costs associated with selling or buying a home, which can include taxation, real estate agent fees, and other costs.

Households that do not own a home must participate in the rental market. I assume that foreign or institutional investors, who are not explicitly modelled, supply housing in the rental market, and I abstract from the

⁴⁹Appendix 3.5.7 shows that results are robust to several specifications where $H = 3$.

⁵⁰Some papers, like Li, Liu, Yang and Yao (2015), distinguish between idiosyncratic and aggregate house price shocks. My main results do not explore that possibility, but Appendix 3.5.6 studies a case in which idiosyncratic house price risk is correlated with labor income risk.

equilibrium determination of house prices for tractability and simplicity⁵¹.

The rental price $r_t^s(\Omega_t^h)$ depends on current housing prices $p_t^h(\Omega_t^h)$:

$$r_t^s(\Omega_t^h) = \gamma^r p_t^h(\Omega_t^h). \quad (37)$$

I assume that the government provides housing aid to income-poor households for whom rental costs are large. In particular, the government pays all rent that is above 30% of household income. This is a stylized representation of housing aid programs in the United States, in particular the Section 8 program (Housing Act of 1937), which provides families with low income with Housing Choice Vouchers or project based assistance. In the PSID data roughly 2% of working age households receive this subsidy. In the model this fraction is about 3%.

During the working period, households are subject to exogenous moving shocks with probability π_{hm} . They represent events such as finding a new job in a different place or suffering a job relocation. In the model, when the moving shock realizes, agents sell their houses at the beginning of period, before they take their consumption and saving decisions. They must then spend that period in rental housing but can freely reoptimize afterwards. This specification captures that, depending on the stages of life and income realizations, agents might optimally choose to rent in the new location even if they were owners before.

3.4.3.6 Mortgages

When a household wants to acquire a house of quality j , it can apply for a loan or mortgage m_t . I define mortgages so that $m_t \leq 0$. In order to get it, the household must fulfill two conditions: a downpayment or loan-to-value (LTV) restriction and an income test or loan-to-income (LTI) restriction.

$$m_{t+1} \geq -\lambda_h p_t^{h_j}(\Omega_t^h) \quad (38)$$

$$m_{t+1} \geq -\lambda_y y_{it}(\Omega_t^y) \quad (39)$$

⁵¹Section 3.6.4 contains an approximation to how my counterfactual results would change under endogenous determination of housing prices

where $\lambda_h < 1$. There is no uncollateralized debt, so households can only get indebted when they buy a house.

Borrowers pay an exogenous interest rate on their debt r^b which is larger than the risk-free rate r^a . Households decide on their repayment schedule, but in every period they must at least pay the interest accrued by their debts and cannot reach their terminal age T with an unpaid mortgage balance, even if their net worth is positive (as in [Attanasio et al. \(2012\)](#)).

$$m_{t+1} \geq \frac{m_t}{1 + r^b} \quad (40)$$

$$m_T = 0 \quad (41)$$

Within this framework, households can extract equity from their homes in two ways. First, they can sell them and either move to rental housing or buy a new smaller or cheaper house. Second, they can decide to delay the repayment of the mortgage principal, thus extending their mortgage duration. For simplicity, I assume that they cannot increase the principal of their debt by remortgaging or accessing home equity lines of credit.

To reduce the dimensionality of the problem, due to computational considerations, I assume that households cannot simultaneously hold both a mortgage m_t , risk-free assets a_t , and risky assets f_t , but only two of the three. This assumption is weaker than modeling mortgages as negative safe assets, because it still allows households in debt to make a choice between positive safe and risky assets, as long as the choice is not interior. Therefore, mortgagors in the model are able to hold liquidity without incurring the participation cost to the stock market.

$$a_{t+1}f_{t+1}m_{t+1} = 0 \quad (42)$$

3.4.4 The government

Disposable income $\lambda(y_{i,t})$ is obtained from pre-tax income $y_{i,t}$ using the tax function $\lambda(\cdot)$ ([Benabou \(2002\)](#), [Heathcote, Storesletten and Violante \(2014\)](#)):

$$\lambda(y_{i,t}) = \lambda y_{i,t}^{1-\tau} \quad (43)$$

This specification can be negative at lower income levels and thus includes, in a parsimonious way, both progressive labor income taxation and many income-tested welfare programs, such as unemployment insurance, EITC, food stamps, etc.

The government also taxes capital income from risky and safe assets at a flat-rate τ_a . It uses the proceedings to finance useless government spending g and social security for old people $p(\cdot)$. The latter is a function of a household's last income realization.

Households can deduct mortgage interest from their labor income tax. Both in the US tax code and in the model, they can choose between getting the standard deduction, which is a fixed amount, and itemization, which implies that they individually deduct qualifying expenses such as mortgage interest. Thus, only households who have a sufficiently large mortgage get the mortgage interest deduction. Furthermore, stock market losses are deductible against asset income and labor income up to \$3,000.

3.4.5 Timing

At the beginning of the period, households learn the common realization of the aggregate state Ω_t , which implies that they find out about housing prices $p_t^h(\Omega_t^h)$ and stock returns $r_t^f(\Omega_t^f)$, and their individual realization of labor income $y_t(\Omega_t^y)$. Jointly, those determine their net worth or cash-on-hand in period t :

$$\begin{aligned} coh_t = & p_t^h(\Omega_t^h)h_t + (1 + r_t^f(\Omega_t^f)(1 - \tau_a))f_t + \\ & (1 + r^a(1 - \tau_a))a_t + (1 + r^b)m_t + T(y_t(\Omega_t^y), m_t) \end{aligned} \quad (44)$$

where $\lambda(\cdot)$ represents progressive taxation of labor earnings net of mortgage interest payments.

Households get utility from their housing stock h_t at the beginning of the period. Then they decide on their consumption c_t and their savings for the next period, which are composed of their liquid accounts a_{t+1} , stocks f_{t+1} , and housing h_{t+1} , minus any outstanding mortgage balance m_{t+1} .

Both in the model and in the data, a household can have negative net

worth. In the model, that is represented by $coh_t < 0$ and can arise when a household suffers a negative housing price, income, or financial shock while holding a significant mortgage. Households can continue to hold their house as long as they are able to make interest payments to their mortgage out of their financial savings or labor income.

If a household has exhausted all of their financial assets, cannot make interest payments to their mortgage, and cannot pay for all of its debt even after selling its house, it goes bankrupt. They return the keys of their house to the bank, their debt is cancelled, and suffer a utility penalty, which incorporates stigma effects and the negative consequences of a bankruptcy flag on future credit reports.

$$\begin{aligned}
& \text{If } coh_t < 0 \\
& \text{and } T(y_t(\Omega_t^y), m_t) + r^b m_t + (1 + r^a(1 - \tau_a))a_t + \\
& \quad (1 + r_t^f(\Omega_t^f)(1 - \tau_a))f_t < 0, \\
& a_{t+1} = 0, f_{t+1} = 0, h_{t+1} = 0, c_t = 0.01, m_{t+1} = 0
\end{aligned} \tag{45}$$

3.4.6 Budget constraint

The period by period budget constraint is:

$$\begin{aligned}
& p_t^h(\Omega_t^h)h_{t+1} + \kappa^h p_t^h(\Omega_t^h)h_{t+1}\mathbb{I}(h_{t+1} \neq h_t) + r_t^s(\Omega_t^h)\mathbb{I}(h_t = 0) + \\
& \quad f_{t+1} + \kappa^f \mathbb{I}(f_{t+1} > 0, f_t = 0) + a_{t+1} + m_{t+1} + c_t = \\
& p_t^h(\Omega_t^h)h_t + (1 + r_t^f(\Omega_t^f)(1 - \tau_a))f_t + (1 + r^a(1 - \tau_a))a_t + \\
& \quad (1 + r^b)m_t + T(y_t(\Omega_t^y), m_t)
\end{aligned} \tag{46}$$

where $T(y, m)$ represents the tax system described in Section 3.4.4.

3.4.7 Household's problem

Working-age households. They solve the following problem:

$$U_t(y, a, h, f, m, \Omega) = \max_{c, a', h', f', m'} \left\{ [(\theta_t c_t^\nu s_t^{1-\nu})^{\frac{(\psi-1)}{\psi}} + \right. \quad (47)$$

$$\left. \beta(\mathbb{E}_t U_{t+1}(y', a', h', f', m', \Omega')^{1-\gamma})^{\frac{1}{1-\gamma} \frac{\psi-1}{\psi}} \right\} \quad (48)$$

subject to the no-shorting condition for safe and risky assets (32, 33), LTV and LTI constraints when buying a home (38 and 39), the requirement to at least pay interest on debt in every period (40), the restriction on holding both risky and safe assets while having a mortgage (42), the bankruptcy condition (45), and the budget constraint (46).

Retired households. Their social security income p is a function of their last realization of labour earnings before mandatory retirement (they cannot retire before 65). They solve the following problem (where y_l is their last realization of income before retirement):

$$U_t(y_l, a, h, f, m, \Omega) = \max_{c, a', f', h', m'} \left\{ [(\theta_t c_t^\nu s_t^{1-\nu})^{\frac{(\psi-1)}{\psi}} + \right. \quad (49)$$

$$\left. \beta \xi_t (\mathbb{E}_t U_{t+1}(y_l, a', h', f', m', \Omega')^{1-\gamma})^{\frac{1}{1-\gamma} \frac{\psi-1}{\psi}} + (1 - \xi_t)v(b)]^{\frac{\psi}{\psi-1}} \right\}$$

where $v(b)$ is determined by Equation 30. Their maximization problem is subject to the no-shorting condition for safe and risky assets (32, 33), LTV and LTI constraints when buying a home (38 and 39), the requirement to at least pay interest on debt in every period (40), the restriction on holding both risky and safe assets while having a mortgage (42), the bankruptcy condition (45), and a budget constraint with no income risk (50).

$$\begin{aligned} p_t^h(\Omega_t^h)h_{t+1} + \kappa^h p_t^h(\Omega_t^h)I_t^h + f_{t+1} + \kappa^f I_t^f + \\ a_{t+1} + m_{t+1} + c_t + r_t^s(\Omega_t^h)\mathbb{I}(h_t = 0) = \\ p^h(\Omega_t^h)h_t + (1 + r_t^f(\Omega_t^f)(1 - \tau_a))f_t + \\ (1 + r^a(1 - \tau_a))a_t + (1 + r^b)m_t + T(p(y_l), m_t) \end{aligned} \quad (50)$$

where $p(\cdot)$ represents social security.

3.5 Calibration

3.5.1 Aggregate state

The aggregate state of the economy in a calendar year Ω_t is the combination of three related elements⁵²: the state of the labor market Ω_t^y , the state of the stock market Ω_t^f , and the state of the housing market Ω_t^h .

There are two possible realizations for the aggregate state of the labor market Ω^y , which correspond to expansions and recessions. In the data, I define a period to be recessionary if any part of it falls under an NBER-defined recession. The state of the labor market in $t + 1$ determines the conditional distribution of shocks that agents face given their earnings in t , as described in Section 3.3.

I discretize the stock market state Ω^f space in four possible realizations. These are obtained by splitting the distribution of yearly stock market returns during my sample period (1963-2015) into terciles. Each of the top three states of Ω^f corresponds to a realization of r^f equal to the average return for each of these three terciles. Additionally, I include a disaster state, that corresponds to the average of the lowest 5% of annual stock market realizations during this period. Taking into account the possibility of a disaster in the stock market is important to understand the low levels of stock market participation and the equity premium puzzle (Bansal and Yaron (2004), Barro (2006)), as well as the age patterns of stockholding (Fagereng et al., 2017). The framework I propose extends these previous studies by letting stock market states Ω^f , including the disaster state, and the aggregate state of the labor market Ω^y be correlated.

The housing aggregate state Ω^h is modeled in a similar fashion, but its memory is longer. Households know the current realization of house prices, and whether they have grown or decreased from the previous period. This can equivalently be understood as two separate states (current house prices

⁵²For the description of the model, t indexed both year and age, which were equivalent from the perspective of a cohort. Naturally, calendar years and their associated states happen at different ages for different cohorts. To keep the notation in this section clear, I describe it from the perspective of a single cohort.

and current house price *growth regime*), or as a restricted Markov 2 process for housing prices, in which $Pr(p_{t+1}^h = x | p_t^h = y, p_{t-1}^h = z)$ is the same for any (y, z) such that $y < z$ and for any (y, z) such that $z > y$. Given that house price growth regimes are persistent in the data, households expect house prices to continue growing when they have grown in the past. Housing prices are discretized to four possible realizations in every period.

In the simulation, the realizations of the aggregate state Ω_t correspond to their counterparts in the data for each specific year. For instance, when agents of the oldest cohort reach 53 years of age they face a good realization of the stock market aggregate state because they were born in 1942 and 1995 was a year of high stock returns. The left hand side panel of Figure 30 shows how biennial stock returns in the model (blue line) closely approximate real returns in the S&P500 (red line). The right hand side panel of Figure 30 shows how the specific implementation of the housing price-to-income (PTI) ratio in the model (blue line) compares with PSID data on the ratio of median house prices to median household labor income in PSID data (red line). Episodes of high house prices and episodes of high house price growth are relatively persistent.

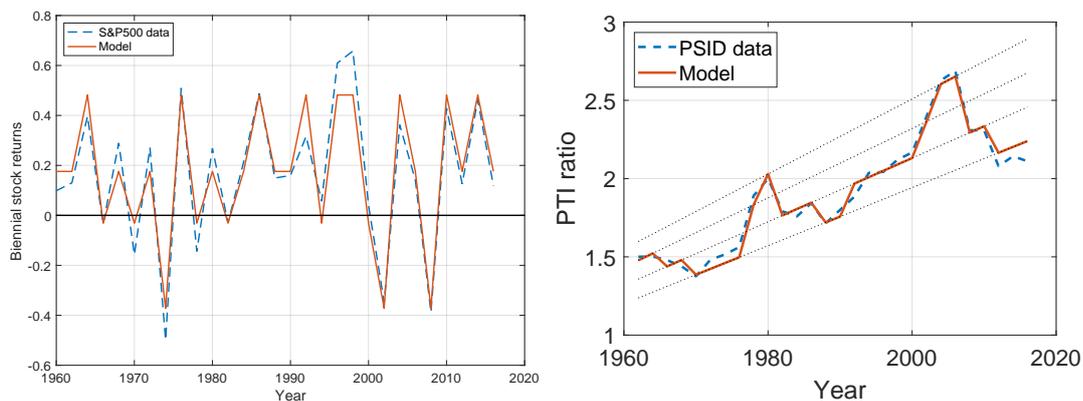


Figure 30: Stock market returns, housing median price-to-income ratio

From the perspective of the agents in the model the realization of the aggregate state is stochastic, so it is necessary to determine how they form predictions over it, and in particular how the correlation between the different elements of the aggregate state is perceived by the agents.

I assume that house price growth regimes and recessions are persistent,

that they might be mutually correlated, and that the conditional distribution of stock market returns depends on whether we are in an expansion or in a recession. The conditional distribution of house price realizations depends then on the current house price and on the current house price growth regime.

Thinking of these assumptions as restrictions on a flexible empirical Markov transition matrix, they imply:⁵³

$$\begin{aligned} Pr(\Omega_{t+1}^h, \Omega_{t+1}^{hg}, \Omega_{t+1}^y, \Omega_{t+1}^f | \Omega_t^h, \Omega_t^{hg}, \Omega_t^y, \Omega_t^f) = \\ Pr(\Omega_{t+1}^f | \Omega_{t+1}^y) Pr(\Omega_{t+1}^h | \Omega_{t+1}^{hg}, \Omega_t^h) Pr(\Omega_{t+1}^{hg}, \Omega_{t+1}^y | \Omega_t^{hg}, \Omega_t^y) \end{aligned} \quad (51)$$

Those dependences are allowed to be very flexible. I empirically estimate from the data the probability of stock returns being disastrous, low, normal, or high in an expansion and a recession, and use it to define $Pr(\Omega_{t+1}^f | \Omega_{t+1}^y)$. $Pr(\Omega_{t+1}^h | \Omega_{t+1}^{hg}, \Omega_t^h)$ is defined according to the average house price growth or decrease associated with each of the two states of Ω^{hg} . I estimate $Pr(\Omega_{t+1}^{hg}, \Omega_{t+1}^y | \Omega_t^{hg}, \Omega_t^y)$ directly from their empirical counterparts to obtain an 4x4 transition matrix.

The restrictions I impose imply assuming that households do not use certain information to make their predictions. First, I assume that agents do not use the state of the housing market Ω_t^h or Ω_t^{hg} to directly predict the realization of the stock market state Ω_t^f . This is empirically justified by the low correlation between the housing state and stock market returns. For instance, a regression of stock returns on the housing state and the recession state yields that the former is insignificant while the latter is significant at the 5 percent level.

Second, I assume that agents do not use the state of the stock market Ω_t^f to predict the aggregate state Ω_{t+1} . This restriction would be violated if, in the data, stock returns in a given year were a strong predictor of stock returns two years later, or of a recession two years later.⁵⁴ How-

⁵³For purposes of this representation Ω^{hg} represents housing growth regimes and Ω^h represents house prices.

⁵⁴A significant part of the literature has established that price-dividend and earnings-price ratios are predictors of future stock returns (Campbell and Yogo, 2006), although some of the relationships between economic and financial variables and future stock performance are unstable and change over time (Pesaran and Timmermann, 1995), and sometimes they react to studies being published about

ever, this assumption does not mean that stock returns are fully i.i.d. in this model, as the correlation between them and the aggregate state in the labor market induces some persistence in high (low) returns during expansions (recessions).

I impose these restrictions due to several considerations. First, with around 50 years of comparable data, the estimation of an empirical transition matrix that allows for all possible correlations would be very noisy. Second, it is not clear to which extent that matrix, even if it could be estimated, would be incorporated in household decision making, in particular taking into account that agents would need to know it *ex ante*. In this context, assuming that households know the persistence of expansions and recessions, of house price growth regimes, and the correlation between stock returns and the aggregate state of the labor market is less stringent than assuming that they know the full correlation structure amongst all possible shocks. In any case, the precise structure of household expectations about the movement of aggregate variables and how it is updated over time remains an open question.

3.5.2 Externally calibrated parameters

Table 5 represents the most relevant externally calibrated parameters and their sources. For presentation purposes, all variables and parameters that correspond to a time period are presented in annual terms, and converted to biennial terms in the model.

I set the risk aversion coefficient to 4. While this is on the higher side of usual estimates in the macro literature, it lies somehow on the lower side of the values used by the finance literature to rationalize the equity premium puzzle in specifications with Epstein-Zin preferences (e.g. [Bansal and Yaron \(2004\)](#) use 10, [Campanale et al. \(2015\)](#) use 5, etc.). The elasticity of intertemporal substitution is more disputed in the literature. In the presence of disaster risk, in models in which asset prices are endogenous, an elasticity above one is needed to make the probability of a disaster and asset prices inversely related ([Barro, 2009](#)). I follow [Kaplan and Violante \(2014\)](#) for its exact quantification (see their footnote 28 for a discussion

them ([McLean and Pontiff, 2016](#)).

regarding this estimate) and set it to 1.5.

I establish the risk-free rate at its historical (real) average of 1%, plus an additional 1% to account for the liquidity services of risk-free money. The mortgage interest rate is at its historical average over the life of this cohort of 4%, and it is 1% higher for retired people to reflect the more stringent credit conditions they are subject to. This assumption is much looser than assuming that retired people cannot get a mortgage or buy a home.

I assume that the downpayment required to get a mortgage is 20% of the value of the house, and that the income test consists in having yearly household income that is at least 1/9th of the value of the mortgage. These are standard in the literature and are roughly the average conditions in the United States during this time period.

For the social security replacement rate, I follow studies that have empirically estimated it from household data. Frequently used values for tax progressivity with the tax function represented in Equation 43 are around 0.15-0.18 (0.151 in [Heathcote et al. \(2014\)](#)). However, in this study I am considering explicitly that US households can choose between deducting a fixed amount from their income tax bill (standard deduction) or deduct a set of qualifying expenses, which in the case of this model is their mortgage interest. These elements strongly affect the progressivity of the system (in particular, the standard deduction makes the tax system more progressive), so directly borrowing those coefficients would result in biased estimates of household tax bills. In order to tackle this issue, I re-estimate the progressivity coefficient from PSID data following the procedure described in Appendix 3.1.1.2. I set the parameter that controls average taxation λ to the level that implies an average tax rate of 35% for the average household, close to the historical level for the 1940s generation comprising federal and state taxes and FICA contributions. With respect to the standard deduction, I set it at a level that implies that the percentage of people choosing to itemize is close to the data, which is around 30% across all ages. This level, 6% of average income, is lower than its historical levels (e.g, around 10% of average income in the early 70s) because the model abstracts from itemizable expenses other than mortgage interest

and local property taxes, such as out-of-pocket medical expenditure, state taxes, charitable contributions, etc.

As for the bankruptcy penalty, I assume that going bankrupt makes households as unhappy as consuming 15% of average income for a period (Equation 45), which keeps bankruptcy rates for the 1940s generation very low. Housing adjustment costs are around 10% of the value of the property (Smith, Rosen and Fallis, 1988), which I distribute equally amongst seller and buyer.

Risk aversion	γ	4	Kaplan and Violante (2014)
EIS	ψ	1.5	Kaplan and Violante (2014)
Housing utility share	ν	0.2	NIPA data
Risk-free interest rate	r^a	2%	
Mortgage interest rate	r^b	4%	
LTV restriction	λ_h	0.8	Downpayment 20%
LTI restriction	λ_y	9	Johnson and Li (2010)
Taxation level	λ	0.64	See text
Progressivity	τ	0.085	See text
Soc. sec. replacement rate	$p(\cdot)$	55%	Mitchell and Phillips (2006)
Housing adjustment cost	κ^h	5%	Smith et al. (1988)
Standard deduction	sd	6%	See text

Table 5: Externally calibrated parameters and sources

3.5.3 Internally calibrated parameters, targets, and model fit

The model has 7 free parameters which are jointly calibrated to match 7 targets in the data. I perform the calibration for the 1940s cohort, and then keep them constant across cohorts in the experiments unless otherwise specified. Table 6 summarizes the data and the parameter which is more closely related with each of the targets.

The wealth to income ratio of 3.1 is standard in macroeconomic studies and corresponds to the wealth to income ratio of the bottom 95% of the wealth distribution, which I am focusing on. I obtain house ownership data from the PSID, stock market participation from the SCF, and bequest

Moment	Data	Model	Key parameter	Value
W/Y ratio	3.1	3.1	Discount factor β	0.930
Avg. bequest (/income)	2.7	2.7	Bequest taste ϕ_1	4.7
% leaving no bequests	20%	16%	Bequest taste ϕ_2	6.4
Homeownership at age 40	77%	78%		
... of detached houses	68%	71%	Housing taste s_2	8.5
... of other housing	9%	7%	Housing taste s_1	2.0
% buying houses at age 40	4.5%	4.2%	Moving shock π_{hm}	0.05
Stock participation, age 40	36 %	36%	Participation cost k^f	0.25

Table 6: Targeted moments, model fit, and calibration

targets from [Hurd and Smith \(2001\)](#), adjusted for this specific cohort (see [Appendix C](#) for more details).

Matching homeownership at a particular age allows me to get an estimate for the extent to which households enjoy living in owner-occupied housing, over and above its value as a financial investment and collateral. On the other hand, getting the level of stock market participation right at a relatively early age allows me to discipline the stock market participation cost κ^f . This parameter is not straightforward to estimate directly from the data, as it not only includes direct costs such as opening a brokerage account, but also the opportunity costs generated by spending time acquiring information about the stock market. Finally, the percentage of people buying houses after prime homebuying age is informative of the number of people who are moving for reasons that I do not model explicitly, which I summarize in the moving shock. Getting this probability right is relevant to appropriately capture that homeowners are sometimes forced to liquidate their houses and move somewhere else, which has associated transaction costs and increases the riskiness of owner-occupied housing as an investment.

As [Table 6](#) shows, the model fits its targets very well with the associated calibrated coefficients. The discount rate is relatively low with respect to what is standard in a one-asset model. Households value housing, and the utility value of owner-occupied houses provides a further motive to

hold assets beyond life-cycle and precautionary savings, which reduces the calibrated value of household patience. Besides, relatively low levels of β are also frequent in the portfolio choice literature when stocks are available as an investment option with high returns.

The housing taste parameters s_2 and s_3 do not have a direct interpretation. The calibration for s_3 implies that, for an agent who is currently consuming the average level of labor income, living in a large owner-occupied house provides, *ceteris paribus*, the same utility increasing nondurable consumption by 70%. There is a 5% yearly probability of receiving a shock that forces the household to move. A particularly relevant parameter is the one-off cost to start participating in the stock market k^f , which is calibrated to be 25% of average yearly earnings.

There is scarce data about the initial wealth of the 1940s cohort at labor market entry. However, I can observe their homeownership and stock market participation rates. Thus, I set the initial condition of the model to the most conservative possibility that is consistent with the observed homeownership (20% equity on the house for the initial homeowners) and stock market participation (1\$ in stocks for the initial stockholders). In Appendix 3.5.1 I provide results for the case in which all agents start at zero wealth. All conclusions are unchanged, although the model with initial zero wealth underestimates homeownership at earlier ages.

Appendix 3.2 briefly describes the solution method of the model.

3.6 Results

3.6.1 Untargeted moments, 1940s cohort

The model replicates life-cycle homeownership profiles and the patterns of house buying by age for the 1940s cohort very well (Figure 31, top panel). Both in the model and in the data, most households become homeowners between ages 20-35, and then the share of households that live in their own home stabilizes around 80%.

It also generates a share of households participating in the stock market that increases like in the data (Figure 31, bottom left panel). Standard portfolio choice models struggle to generate the low levels of participatio

observed amongst the young. In the model, young households do not participate in the stock market because they are concentrating their resources in saving for a downpayment and starting to pay their mortgages, rather than spending time and resources in acquiring information and access to the stock market.

Many households hold mortgages at the same time as they start investing in stocks. The model's flexibility implies that it replicates this fact, which is relevant for household wealth accumulation and to replicate household liquidity positions. Figure 31, bottom right panel, shows that in the model households pay back their mortgages slowly, a feature which is not targeted in the calibration nor in the model definition of mortgage, which does not include products like fixed 30-year loans. Thus, the model suggests that the horizon of available mortgage products closely resembles what households would choose if they were to freely decide on their repayment schedule.

The model is also successful in replicating portfolio patterns by wealth (Figure 32 reports them for the retirement age to avoid confounding life-cycle effects in the wealth distribution). As stressed in e.g. [Gomes and Michaelides \(2005\)](#), a standard portfolio choice model would yield stock holding patterns which are mildly *decreasing* rather than increasing in wealth, while the latter is true in the data. In this model, households of low to middle income buy houses and concentrate a large part of their resources in their residential investments instead of acquiring information and access to the risky stock market, which helps to explain why their participation rate is relatively low. Besides, the correlation between labor income risk and stock market returns further reduces their incentives to participate in the stock market. Richer individuals, on the other hand, have sufficient resources available even after buying their homes, and they invest them in the stock market, in which they reap higher returns that in turn make them wealthier. They still choose to keep some liquidity, but, like in the data, it is a very small proportion of their total amount of assets.

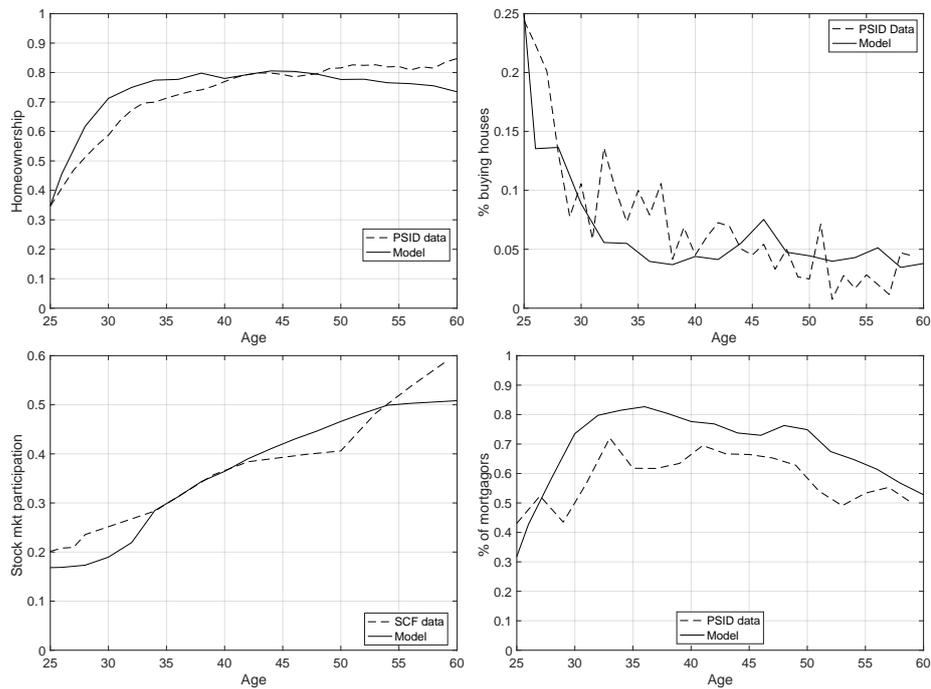


Figure 31: Life-cycle profiles for the 1940s cohort. Top left: homeownership by age; top right: proportion of households buying a house by age; bottom left: stock market participation; bottom right: percentage of all households with a mortgage by age.

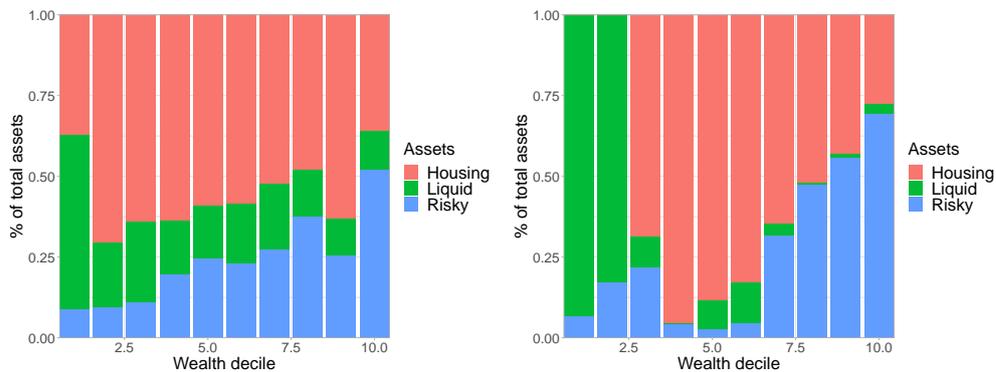


Figure 32: Bottom: portfolio shares of assets by wealth decile at retirement age (left: PSID data, right: model).

3.6.2 Explaining intergenerational differences in homeownership

Keeping constant the preference parameters that I have calibrated to the 1940s cohort, I now turn to studying which are the key intergenerational changes that explain the reduction in homeownership for younger cohorts.

In this experiment, cohorts differ in four ways. First, younger cohorts face more unequal and riskier earnings processes, as described in Section 3.2.1. Second, the exogenous house prices and stock returns correspond to those that each generation actually faced, thus implying that, for younger generations, the median earner needs to spend more years of income to buy a house. Third, there have been changes in financial conditions. On the one hand, different mortgage products were available to the 1980s generation during their homebuying years, which I replicate as a reduction in downpayment requirements. Namely, I assume that the maximum LTV ratios of mortgages increased from their baseline level of 80% to 100% between 2000 and 2010, after which they unexpectedly went back to normal.⁵⁵ On the other hand, I reduce stock market participation costs to match the stock market participation profile (see Section 3.6.3). Fourth, I input to each generation their specific average demographic profile by age, which captures the effect on consumption needs of differential timings in marriage and childbearing.⁵⁶ For a cleaner comparison, the initial condition that captures the percentage of households that enter the model as homeowners does not change across generations.⁵⁷

Figure 33 shows the homeownership rates for each of the three cohorts in the data, compared with the profile implied by the model. The line for the 1940s replicates what I have shown in Figure 31. Notably, keeping preference parameters and mortgage conditions constant, the model very closely replicates the decrease in homeownership that occurred between

⁵⁵Duca, Muellbauer and Murphy (2011), using American Housing Survey data, show that average LTV ratios for first time buyers, which were stable around 0.80-0.85 in the 1980s and early 1990s, jumped up to 0.90-0.95 during the 2000s. Glaeser, Gottlieb and Gyourko (2012) use housing industry data and show that for most of the 1998-2008 period the 75th percentile of LTV ratios at origination was above 95%, with the 90th percentile consistently around 100%.

⁵⁶Appendix 3.3.7 represents these OECD equivalence scales, obtained from PSID data, for each of these generations

⁵⁷This assumption is conservative, as it is likely that this percentage has been decreasing over time, as labor market entry takes place later for younger generations.

the 1940s and the 1960s cohort.

Once the transformations in the financial sector are taken into account, the model is successful in explaining the homeownership profile of the 1980s cohort (dotted blue line in Figure 33). When they are ignored, the model generates later homebuying decisions and lower homeownership rates for the 1980s cohort. This difference suggests that the changes in financial conditions were key to prevent homeownership rates of younger cohorts to plummet in a context of unstable, unequal earnings and high house prices.

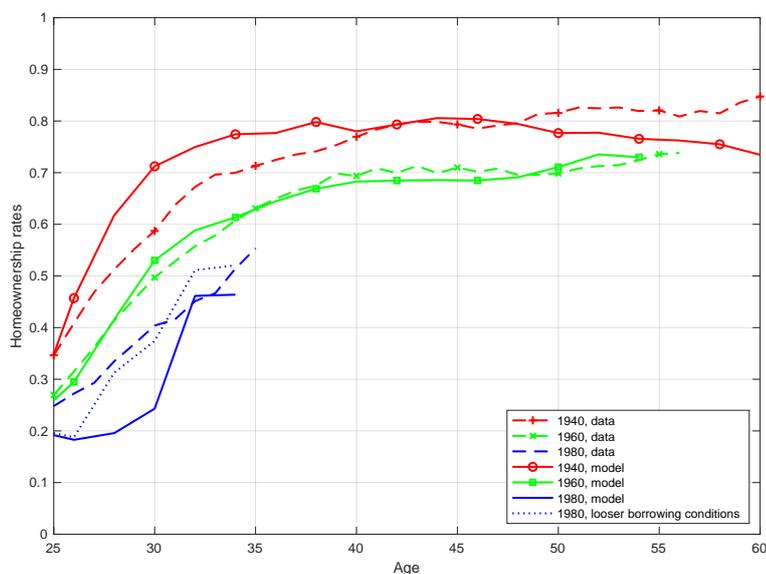


Figure 33: Homeownership by cohorts, data vs. model

3.6.2.1 Decomposing the decrease in homeownership

I now turn to evaluating, using the model, which are the key factors that drove the decrease in homeownership. Table 7 shows the results of a Shapley-Owen decomposition in which I evaluate the relative contributions of six key elements in explaining the reduction in homeownership at different ages: initial earnings inequality, earnings risk thereafter, changes in average housing price-to-income ratios, histories of aggregate shocks, average demographic structure at each age, costs of participation in the stock market and, for the 1980s generation, changes in financial conditions. Given that these are the only differences across cohorts in the model, changing all of them to their corresponding values for the 1940s

generation would imply replicating the model-implied profile for the 1940s. Thus, by counterfactually changing them one by one I can quantify their relative contribution to the difference between the observed profile for a given generation and that of the 1940s.⁵⁸

	1960s generation			1980s generation	
Age	30	40	50	30	35
Total	-9	-8	-9	-14	-22
Earnings	68	48	15	73	38
initial inequality	61	25	-17	41	14
risk	7	23	32	32	24
Aggregates	33	45	91	90	74
house price trend	63	79	46	45	47
histories	-30	-34	45	45	27
Financial conditions	1	-2	-3	-63	-15
stock participation costs	1	-2	-3	5	0
borrowing conditions	0	0	0	-68	-15
Demographics	-2	9	-3	0	3

Table 7: Contribution of each factor in the change in homeownership with respect to the 1940s generation (% of the change), by age

Changes in earnings dynamics are a key driver of the decrease in homeownership rates, although the magnitude of their contribution varies by ages and generations. At age 30, changes in labor market outcomes explain 68 percent of the homeownership gap of the 1960s generation with respect to that born in the 1940s, mostly due to initial earnings inequality. With a more unequal earnings distribution, and little average increases in earnings, households in low ranks of the income distribution have lower initial and expected lifetime earnings than their counterparts in older generations. These households face two issues when they decide whether to buy a house. First, they are financially constrained, as they need to save

⁵⁸All elements in the decomposition have potential interaction effects, which means that shutting them on and off alternatively would not sum to 100% of the changes observed. The Shapley-Owen decomposition allows to obtain the total contribution of each element to the change by considering its contribution to every possible permutation of the other factors being on and off, and averaging over all of these.

for a downpayment and pass an income test. Second, they are aware that having a large mortgage with respect to their incomes is risky, as negative shocks could take them to a situation in which they must reduce a lot their nondurable consumption to make mortgage payments. Thus, in a period of relatively low rental prices, they choose to be renters. For some this is a delay in the decision to buy houses, but for some this state is relatively persistent. At age 40, earnings dynamics still explain almost half of the homeownership gap between generations.

Earnings inequality and risk are closely linked. Even with a constant variance of shocks, higher inequality in the earnings distribution implies that shocks affect people differently. However, to get an intuitive idea of the role of risk, I also check the contribution of changes in earnings dynamics over and above initial realizations. At age 40, at constant initial inequality, riskier earnings explain 23% of the drop in homeownership rates. The higher volatility of earnings discourages households from engaging in a large, risky expenditure like a house. At later ages, the dependence on initial earnings realizations progressively dies out and it is harder to disentangle the effects of initial inequality and risk.

The intuition about earnings inequality and earnings risk is supported by the empirical evidence shown in Figure 34. The gap in homeownership rates between the 1940s and 1960s generations is larger for the lowest earners, which is consistent with the contribution of earnings inequality, but there are also differences all across the earnings distribution, which is consistent with the role of earnings risk.

House prices have increasing relevance to explain the decrease in homeownership as the 1960s generation ages. Initially, this generation entered the labor market in a period of cyclically low house prices, which explains the negative contribution to homeownership of aggregate shocks. In the absence of all other factors, the model predicts that homeownership rates for those born in the 1960s given their histories should be larger than that of the 1940s until age 40. Then the 1960s generation lived through the 2000s boom in house prices. While at the beginning this might have encouraged them to buy houses as an investment, eventually their cost was too large and some households in this generation decided to either wait

until prices decreased or stay as renters.

Despite later household formation and a lower number of children for younger generations, the change in the average number of people in a given household at each age (θ_t in the model), which affects consumption needs, has a very small effect on homeownership rates.⁵⁹ The same applies to changes in stock market participation costs, which I describe in more detail in Section 3.6.3.

The 1980s generation entered the labor market in a radically different period. House prices were high both from a secular and cyclical perspective, but financial constraints were laxer. Prices alone would have explained almost all (90%) of the drop in homeownership at age 30, but the lower downpayment requirements counteracted two thirds of the potential decrease.

The remainder of the difference, over 70 percent, is accounted for by earnings dynamics. For this generation, earnings risk is more relevant than for the 1960s group, which is consistent with the empirical observations in Figure 23 and also with Figure 34, which shows that there was a decrease in homeownership also for the relatively higher earners in this generation.

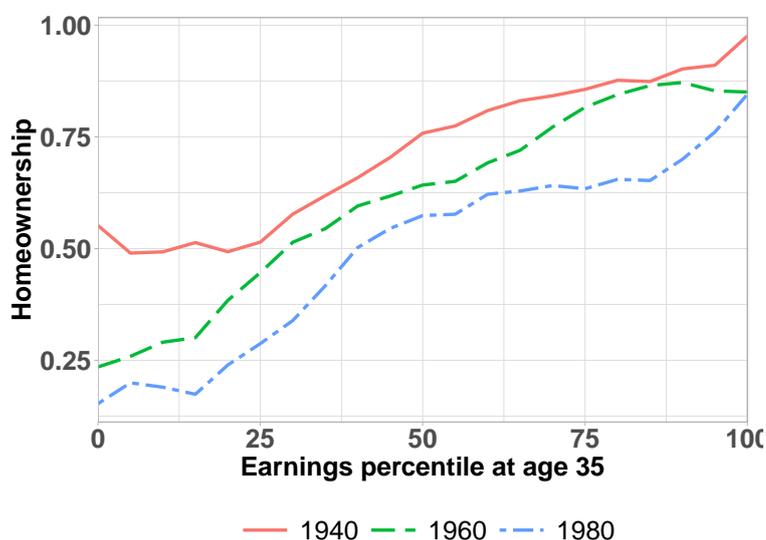


Figure 34: Homeownership by cohorts, by percentile of the earnings distribution at age 35

To more precisely understand the role of earnings dynamics in explain-

⁵⁹This study abstracts from cross-sectional heterogeneity in marital status and fertility choices, which would be an interesting avenue for future research.

ing lower homeownership rates, Figure 35 represents two counterfactual experiments. First (dashed lines), it shows what happens if we attribute the earnings process of the 1940s generation to the younger generations, whilst keeping all else constant. Second (dotted lines), it shows what happens if we keep the initial distribution as observed for each of the generations, but attribute to all of them the earnings *risk* associated with the 1940s cohort. These experiments are different from the previous decomposition because they also take into account possible interaction effects between factors.

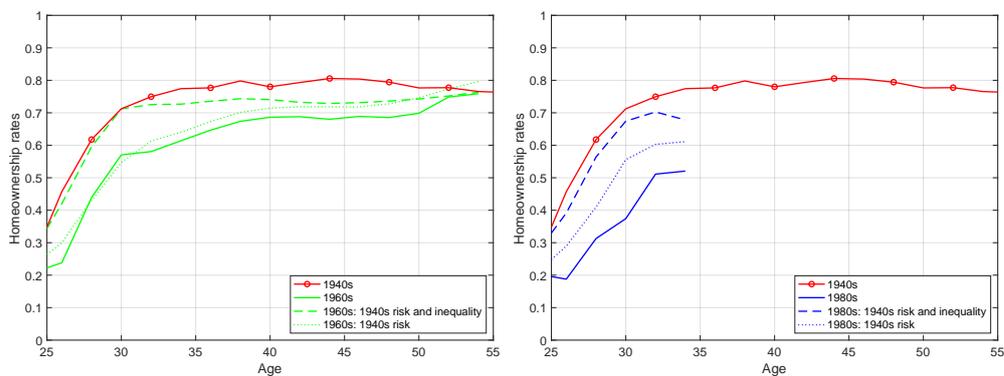


Figure 35: Homeownership by cohorts, the role of earnings risk vs earnings inequality. Left: 1960s generation; right: 1980s generation.

For the 1960s cohort, attributing them the 1940s earnings process would go a long way in closing the homeownership gap between them and the earlier 1940s cohort. For the 1980s generation, endowing them with the earnings inequality and risk that correspond to the 1940s cohort would imply a homeownership rate more than 10 percentage points larger just before age 35.

The results of the risk counterfactual (dotted lines) are close to those of the baseline model, particularly for the 1960s generation. Thus, the reduction in homeownership for this generation is mostly related to the lower earnings of people in lower ranks of the distribution, rather than earnings risk. The role of the latter is more relevant for the 1980s cohort.

In all of these experiments, earnings dynamics are computed on household income, so they embed other factors that changed over the generations such as the timing of family formation. However, these results are

robust to focusing on married couples alone.⁶⁰

Additionally, these counterfactual experiments assume that housing supply is perfectly elastic and so house prices would not react to the increase in housing demand induced by the change in the earnings process. In Section 3.6.4 I relax this assumption and show that a reduction in earnings inequality and risk would imply a significant increase in homeownership for younger cohorts even if we assume that the increase in demand would drive prices up.

3.6.3 Explaining the changes in stock market participation

Understanding the increase in stock market participation documented in Section 3.2.5 requires taking into account not only the changes in earnings dynamics and asset returns, but also the progressive reduction in the cost of access to financial markets over time, which is partially related with the introduction of tax-advantaged, employer-sponsored retirement plans.⁶¹

Figure 36 shows the implications of the model in terms of stock market participation when these changes are taken into account. More specifically, it assumes that stock market participation costs are 30% lower for the 1960s and 70% lower for the 1980s generation than they were for the 1940s generation, and additionally that the initial share of people with positive participation in the stock market has increased over the generations from just below 20% to 25% and 30%. Both of these changes capture the reduction in information costs and the effect of auto-enrolment.⁶²

If the reduction of stock market participation costs is not taken into account, even under the assumption that the initial condition has changed, the profiles generated by the model are counterfactual (central panel of Figure 36). Indeed, the model would predict a reduction rather than an increase of stock market participation over the generations.

The fiscal incentives of IRAs and 401(k) are also not sufficient to explain the increase in stock market participation (right panel of Figure 36).

To gauge this explanation, I modify the nature of the financial asset or

⁶⁰See Appendix 3.5.4.

⁶¹See Section 3.2.4 for details on these and the timing of their introduction.

⁶²Appendix 3.5.5 shows that changing the fixed cost of participation for per-period participation costs can generate similar patterns.

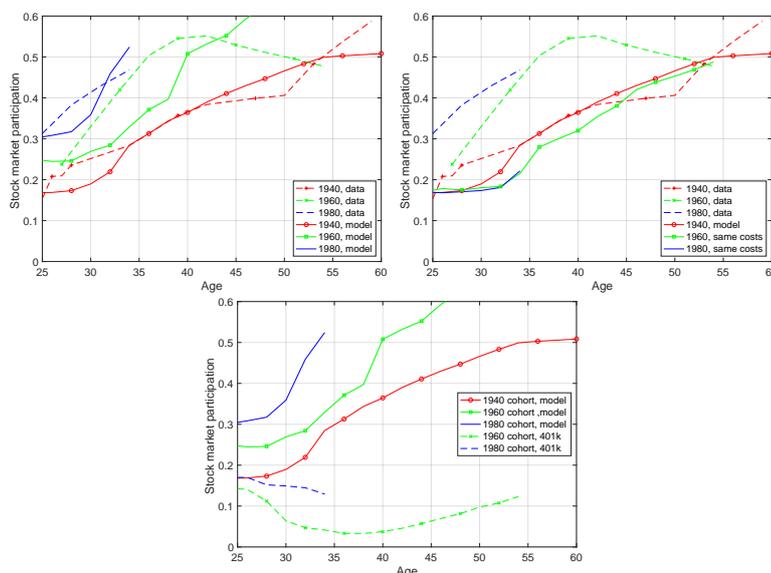


Figure 36: Stock market participation by age and cohort, data vs model. Top left: main model; top right: constant participation costs; right: constant participation costs, stocks with 401(k) tax properties

stock f_t in the model to closely replicate a 401(k). I keep participation costs constant across generations, but assume that contributions to the account are tax-exempt below a certain limit, the interest it generates is tax free, households pay income tax on all amounts withdrawn, and there are penalties for withdrawal before age 60 (10%). Households in the model would react negatively to these features and invest less on the financial asset because of its illiquidity, which makes it costly for households to withdraw from their stocks in response to a bad labor income shock.

3.6.4 Adjusting housing prices

The counterfactual experiments presented so far abstract from general equilibrium effects. However, as the earnings process for the 1960s and 1980s cohort counterfactually changes, so do household decisions, which may impact the evolution of aggregate prices in the economy. It is likely that the increase in housing demand would have had equilibrium effects manifested in an increase in house prices, which could dampen the increase in homeownership rates implied by the experiments.

As such, all results so far can be seen as an upper bound of the possible effects of income dynamics on homeownership, calculated under two equivalent assumptions: either housing supply is perfectly elastic or a non-

modelled investor owns all rental housing and is willing to sell or buy any of it at the observed prices.

In this section, I provide an approximation to these equilibrium effects. I assume that housing supply can be summarized by an isoelastic supply function with elasticity of 1.75, an empirical value estimated for the average U.S. metropolitan area by [Saiz \(2010\)](#). Then I compute the variation in housing prices induced by the increase in housing demand, and find homeownership rates for each cohort under those new prices.⁶³

Figure 37 compares the homeownership rates by age and cohort between the baseline (solid lines), the counterfactual with fully elastic housing supply (dashed lines), and the counterfactual with empirically determined housing supply elasticity (dash-dot lines). Although naturally homeownership is a bit lower for most cohorts and ages, the joint effect of earnings inequality and risk is still very relevant.

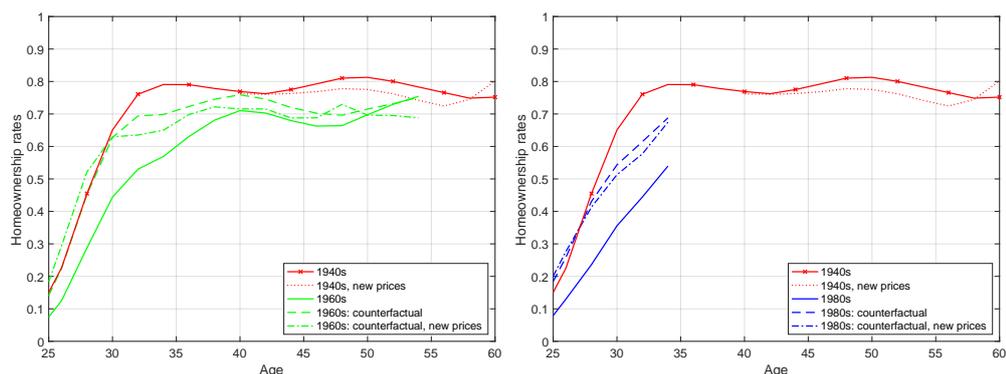


Figure 37: Homeownership by cohorts, benchmark vs. counterfactual earnings processes, empirical housing supply elasticity. Left: 1960s generation, right: 1980s generation.

Therefore, both experiments show that the effect of income dynamics in homeownership is large and can explain a significant amount of the changes in homeownership over the generations, even if we allow for adjustments in average house prices.⁶⁴

⁶³Throughout these experiments, from the perspective of households, house prices are still exogenous shocks. Appendix 3.4.1 contains more details on how the computations in this section are performed and also shows results with fully inelastic housing supply. In this section, I assume households are born with zero wealth.

⁶⁴A more detailed approach could imply modelling the housing supply and rental market sectors, together with a realistic representation of housing devaluation and renovation, and thus obtaining a more flexible formulation of the housing supply function. While such a study could shed light on slow-moving dynamics of housing prices, it is beyond the scope of the model presented in this Chapter given current computational constraints.

All results in this section still assume that stock market returns and income dynamics are exogenous. The stock market is very integrated internationally, so it is less harmful to assume that the US behaves like a small open economy in terms of stock market returns and thus changes in household demand for stocks do not impact stock returns. With respect to income dynamics, the model is already capturing very well their changes over the cohorts and the business cycle, and endogeneizing them would imply losing much of the data-driven richness that is key for the mechanisms considered. However, both are interesting questions that are left open for future research.

3.7 Implications

3.7.1 Reaction to shocks: consumption and homebuying decisions

3.7.1.1 Generational changes

These secular changes in earnings dynamics, wealth accumulation, and portfolio composition have impacted the way households react to shocks. As a simple representation of these, Figure 38 (top left panel) plots [Blundell et al. \(2008\)](#) (BPP) *insurance coefficients*, which represent the percentage of a given shock to persistent earnings that does not get translated into nondurable consumption. The larger this coefficient, the lower the pass-through of earnings shocks to consumption. The decreasing value of this coefficient suggests that younger generations are less insured against income shocks and thus display larger consumption responses when these shocks hit. This is closely related to the reduction in average wealth accumulation for younger generations (Figure 38, top right panel).

The bottom left panel represents the average marginal propensity to consume. in the model, by age and cohort, as a response to a one-off marginally small wealth shock. We observe a strong life-cycle pattern, where MPCs are larger for the youngest, who have lower amounts of wealth to use for self-insurance purposes than the old. These have also increased slightly over generations at earlier ages, suggesting that they are more reactive to shocks.

However, the fact that younger generations own less housing and thus less illiquid wealth has helped to mitigate this effect. The bottom right panel of Figure 38 shows that, if we attribute to younger generations the portfolio of older generations, their MPCs would be even higher.

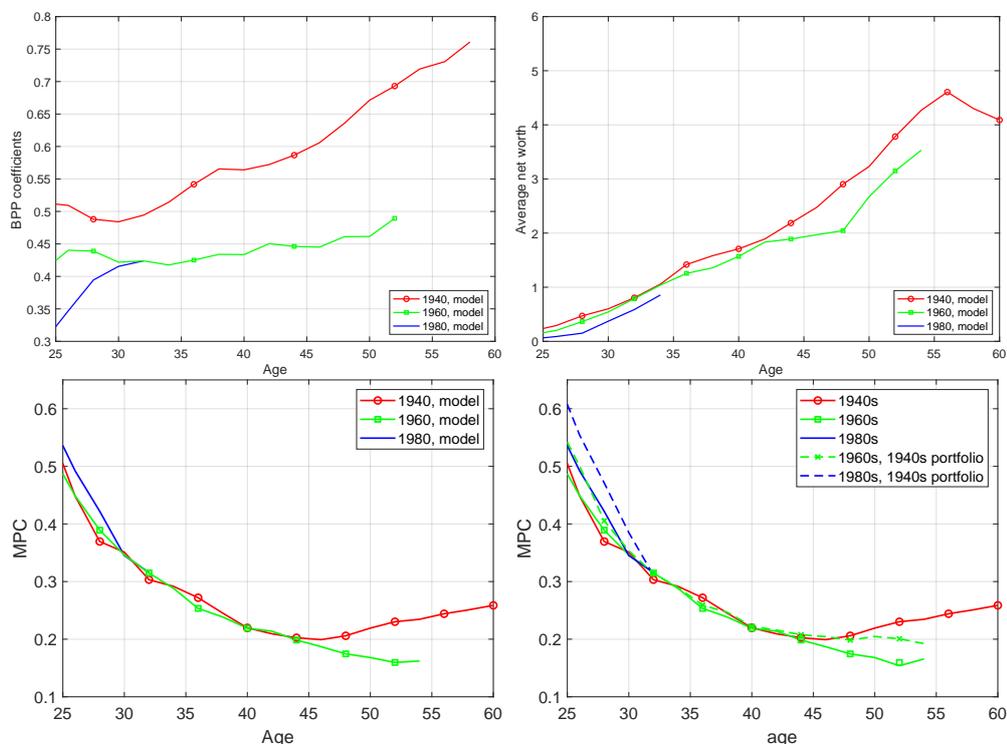


Figure 38: Household responses to shocks. Top left: [Blundell et al. \(2008\)](#) coefficients. Top right: average net worth. Top bottom: average marginal propensity to consume, by age and cohort. Bottom right: including case in which household portfolios are fixed to their 1940s value.

3.7.1.2 Shocks to the aggregate state and the Great Recession

The Great Recession affected both the labor market, the housing market, and the stock market. The model considers explicitly both the histories and the portfolio positions of the different generations that formed the cross-sectional distribution of the economy in 2008, and as such it can shed light on the relative importance of these shocks in explaining household reactions in terms of consumption and homeownership. Furthermore, it provides a measure of how negative housing price expectations need to be in order to generate a large decrease in homeownership as a response to a large decrease in house prices such as that occurred in 2008.⁶⁵

⁶⁵According to the US Census Bureau, homeownership dropped by around 1.5% between its peak in 2006 and two years later. While their definition of homeownership is different from mine (see

Figure 39 shows how an economy that replicates that of the US in 2006 reacts in terms of consumption and housing demand to, respectively, a recession that affects the labor market (red line), a recession with a housing price decrease of 13.5% in two years, which is approximately the observed value for the US between 2006 and 2008 (yellow line), and a recession with a housing price decrease and a negative shock to house price beliefs about the future (purple line). Namely, households, upon seeing the negative shock, expect house prices to start decreasing at 1.8% per year until the negative growth regime is reversed. This is the average decrease that households expect in the model for negative house price growth regimes.⁶⁶ In all cases, the counterfactuals are computed with respect to the case in which none of these events had happened.

Two main messages are apparent. First, looking at the left panel of Figure 39, we observe that the shock to house prices substantially contributed to the reduction in consumption by more than 1.5 percentage points with respect to a standard recession, leading to a decrease in the region of 2-2.5%, depending on the assumptions about housing price beliefs (purple and yellow lines). In the data, consumption dropped 3.4% from peak to trough (De Nardi, French and Benson, 2011).

Second, despite the decrease in homeownership associated with a recession (red line), households react to the decrease in house prices by buying housing because it has become relatively cheap from an intertemporal perspective. An expected persistent decrease of 1.8% a year in house prices reduces the incentives to buy housing, but still generates a positive response of homeownership.

In order to reconcile the dynamics of both homeownership and consumption that we observe in the data with the model, it is necessary to take into account that perspectives on the housing market were particularly bleak during the Great Recession. Namely, I study the case in which households are hit, in 2006, by an unexpected belief shock that implies that they expect house prices to start decreasing by 3.5% each year now and in every future negative house price growth regime (green line).⁶⁷

Section 3.2.5), I use it as a reference for this section.

⁶⁶Appendix 3.4.2 contains more details about how this experiment is performed.

⁶⁷The model is biennial, so this timing assumption approximates the fact that house price expect-

With these assumption on housing beliefs, which is relatively conservative with respect to the actual decrease observed during the Great Recession, the model generates a drop in homeownership and average consumption that broadly aligns with what we observe in the data.

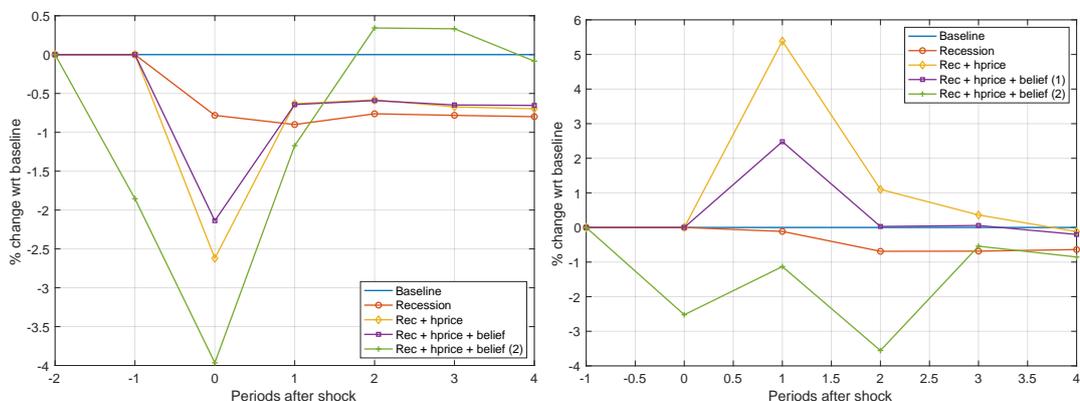


Figure 39: Effects of shocks: left, consumption; right: homeownership.

3.7.2 The future of the 1980s generation

These changes in earnings dynamics and asset returns impact the wealth distribution across generations and social mobility, and can have long-lasting effects that affect the life experiences of the youngest cohorts when compared with older ones. I now turn to using the model to predict, under several assumptions for the evolution of the exogenous earnings process and house prices, the evolution of homeownership rates and other economic variables for the 1980s generation beyond 2015-2020.

In my main experiment (Figure 40, left panel) I assume that house prices will continue to grow on average with respect to median income at a rate of 1.5% per year, and that the earnings process of the 1980s cohort will be that of the 1960s or 1940s cohort for all unobserved years. I then simulate 1000 possible histories of the realizations of house price shocks and aggregate states, consistently with their conditional probabilities in the sample I observe, and plot the median homeownership rates. The dotted lines are percentiles 2.5 and 97.5 within the simulations.

The right hand side panel of Figure 40 shows alternative scenarios.

The black solid line represents the case in which the earnings process of

tations and house prices evolve continuously rather than in a discrete manner. A quarterly version of this model could relax this assumption and generate smoother homeownership profiles.

the 1980s cohort, rather than reverting back to that of the 1940s and 1960s cohort, stays constant from age 35 onwards. Finally, the black dot-dash line represents the case in which house prices are assumed to become constant with certainty with respect to median income from age 35 onwards. For clarity, confidence bands are not reported. Broadly, the median realization of all of these scenarios implies that households born in the 1980s take longer to buy houses, but eventually reach the homeownership rate of the 1960s generation. However, in a large set of scenarios less than 70% of households own houses at age 60.

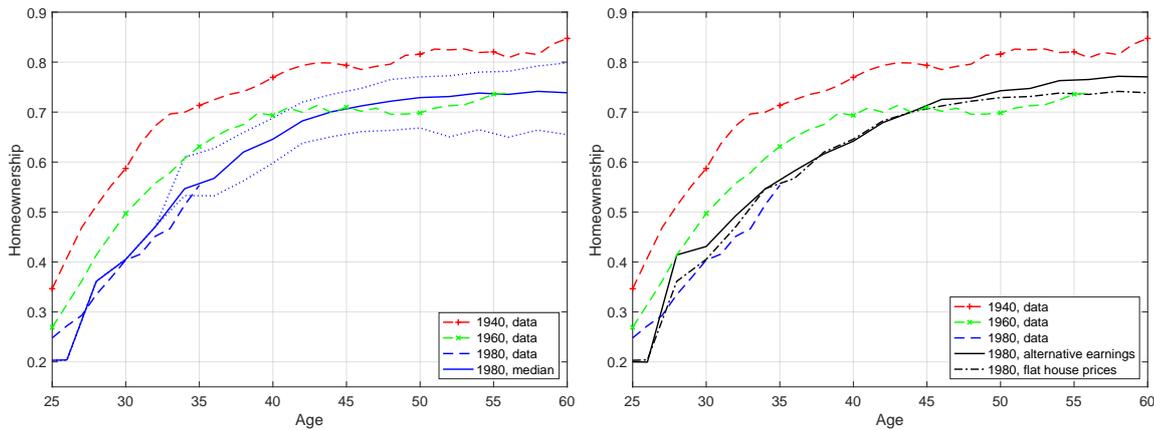


Figure 40: Homeownership: projecting the 1980s cohort into the future

Figure 41 shows associated wealth accumulation profiles for the main experiment. The model predicts slower wealth accumulation for the 1980s cohort, although they eventually reach the 1940s cohort in terms of housing wealth, and most notably that for this generation financial wealth will become an increasing part of their portfolios, particularly if I consider that the cost of access to financial markets is now lower.⁶⁸

However, these averages mask distributional effects. Figure 42 compares the 1960 and 1980s generations by their wealth positions around age 35, and shows that for much of the wealth distribution the increase in stock market participation has not compensated the lower accumulation of housing wealth. Besides, the lower level of housing assets (right panel) shows that the 1980s generation can miss out on potential house appreciations. Thus, the substitution of financial wealth for housing wealth can

⁶⁸These results abstract from possible general equilibrium effects on stock returns induced by the increased accumulation of financial wealth. Appendix 3.5.3 shows that a reduction of 2% in average stock returns would still imply an accumulation of financial wealth comparable to the 1940s cohort.

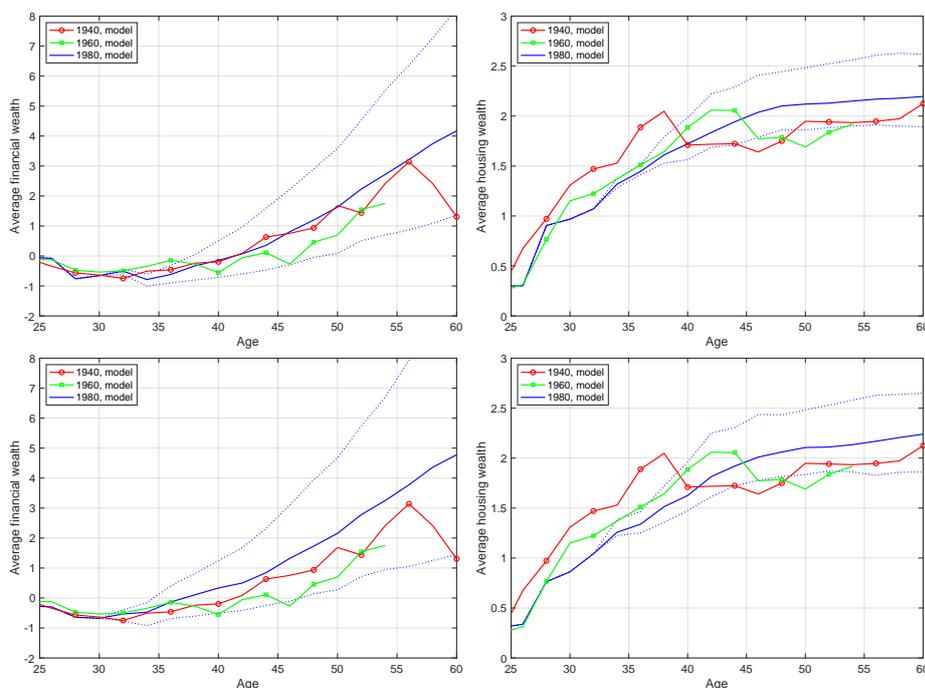


Figure 41: Wealth accumulation: projecting the 1980s cohort into the future. Top: constant stock market participation costs across cohorts; bottom: reduction in stock market participation costs for 1980s cohort. Units are multiples of average income.

be related with an increase in wealth inequality. Table 8 shows the wealth Gini index at retirement for the two older generations and for two different simulations for the 1980s cohort, which differ only in stock market participation costs k^f . The model generates an increase in wealth inequality between the 1940s and 1960s cohort, as we observe in the data, and predicts that, unless stock market participation costs are substantially reduced, wealth inequality will continue to grow. This increase in inequality happens because less households accumulate housing wealth and financial assets are still concentrated amongst the rich. The model predicts that, for wealth inequality not to grow for the 1980s cohort, stock market participation costs need to be reduced such that more than three quarters of the population have access to the stock market by age 60. The larger the share of households that participate in the stock market, the stronger the negative effect on wealth inequality.

In any case, the model underestimates total wealth inequality, as it does not include a set of elements that are important to explain it, such as entrepreneurship or intergenerational links (De Nardi and Fella, 2017). The extent to which these vary over the generations will also have a de-

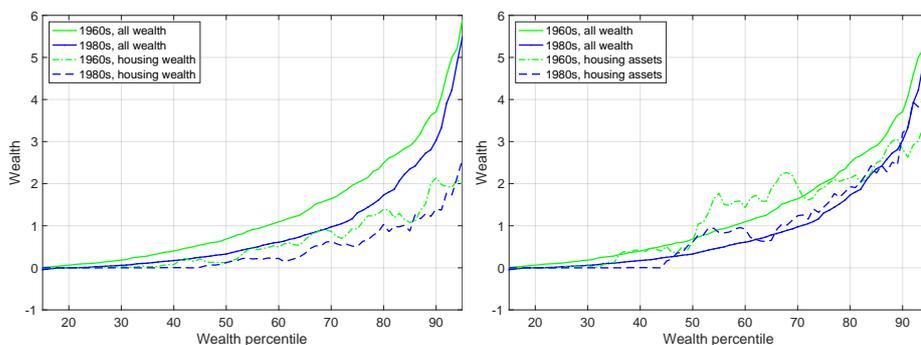


Figure 42: Net worth by wealth percentile, ages 30-40, by generation. SCF data. Units are multiples of average income. For clarity, the top 5% and bottom 15% of the wealth distribution are not reported.

Generation	1940	1960	1980	
Reduction in k^f	-	-	0%	70%
Implied participation at 60	53%	56%	59%	87%
Wealth Gini, data	0.78	0.83	-	-
Wealth Gini, model	0.50	0.53	0.57 (0.026)	0.52 (0.017)

Table 8: Wealth Gini at retirement, data vs model. The data is obtained from people aged 55-64 in the SCF (2001 and 2016 SCF for the 1940s and 1960s cohort respectively). For the model, standard errors from simulation are in parentheses.

terminant effect on the evolution of wealth inequality.

3.8 Extensions and alternative specifications

3.8.1 Canonical earnings process

I now turn to evaluating to which extent a *canonical* earnings process, such as that described by Equations 26-28 in Section 3.3, with changes over generations, can explain the observed changes in homeownership. I obtain its parameters by fitting the cohort-conditional profiles of variances and autocovariances over the life-cycle.⁶⁹

Based on these estimated processes, I recalibrate the model for the 1940s cohort following the same procedure as in my main results, and then simulate the model for the 1960s and 1980s cohort changing only

⁶⁹Appendix 3.4.3 shows the estimated parameters for the canonical process. For this experiment I assume that households start life with zero wealth.

earnings process, asset returns, and financial conditions. Figure 43 shows that the model also generates a decrease in homeownership rates across cohorts, which supports the argument that changes in earnings inequality and earnings risk are key drivers of the observed changes. However, the canonical process substantially overestimates this intergenerational decrease. Earnings are estimated to be a random walk or close to it, and thus the effects of large initial inequality are more persistent than in the more flexible process proposed in this Chapter.

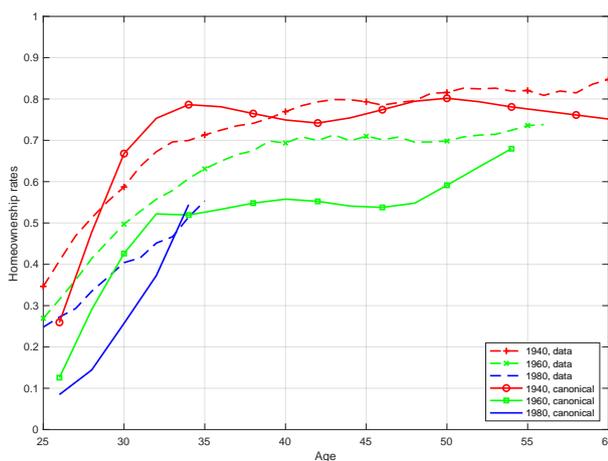


Figure 43: Homeownership by cohorts, data vs. model, canonical earnings process

Additionally, as argued in Section 3.3.1.2, this process has a set of counterfactual implications. It does not replicate the countercyclical skewness of earnings and it implies increases of average earnings when recessions hit. Besides, as shown in recent literature (Gálvez, 2017), it also generates counterfactual implications for stock market participation decisions. Appendix 3.4.3 shows that, in the case of this model, the canonical process generates a counterfactually steeper profile for stock market participation.

3.8.2 Other extensions

Appendix 3.5 shows a set of robustness checks. They show that starting households at zero wealth or at an empirical level of initial wealth, considering changes in marital dynamics and family formation, per-period stock market participation costs, local correlation of income shocks and house prices, and alternative versions of the discretization of houses do not affect the main messages in this Chapter.

3.9 Conclusion

In this Chapter, I study how changes in earnings dynamics over different cohorts have affected their homeownership and portfolio choice decisions. First, I provide empirical evidence, extracted from PSID and SCF data dating back to the 1960s, that there has been a secular increase in household earnings inequality and risk, together with substantial reductions in homeownership and an increase in stock market participation.

Second, I design a flexible earnings process with a business cycle component that accomodates rich, non-normal, non-linear features of earnings risk whilst allowing it to be correlated with the aggregate performance of the economy and asset returns. I find that this process replicates important features of earnings dynamics and their variations over the age and earnings distribution, including the sluggish recovery after a recession.

Third, I develop a rich life-cycle model of housing and portfolio choice that is able to generate key life-cycle and cross-sectional patterns with a relatively parsimonious parametrization. Key elements are a taste for owner-occupied housing, a minimum size for houses, transaction costs, and stock market entry participation costs. I use the model to explain the intergenerational changes I observe in the data without the need of assuming preference changes across generations. Differences in earnings dynamics account for more than half of the reduction in homeownership at ages 30-35.

Looking at the broader implications of my findings, they suggest that taking into account intergenerational changes is important to study household earnings, consumption, and wealth accumulation. At any point in time, the cross-sectional distribution of the economy is formed by many different households who have lived through different histories of shocks at different points in their lives. Acknowledging this fact matters to understand the economic decisions that have led them to be where they are today, and thus to infer parameters to study the effects of policies or the evolution of the economy.

Besides, these results are of interest to policymakers who care about homeownership, intergenerational redistribution, and the evolution of in-

equality. For instance, the model suggests that reducing the costs of access to financial markets can spur wealth accumulation for middle income households, thus reducing overall wealth concentration.

Finally, this Chapter also adds to a burgeoning literature that, based on elements from household finance, points out that considering household portfolio compositions is important to answer many macroeconomic questions, such as consumption responses to shocks or wealth accumulation over the life cycle.

Conclusions

This thesis investigates household earnings dynamics in the United States and in the Netherlands, their changes over time, and their impact on consumption, homeownership, and household portfolio composition. Each chapter has its own conclusions, but all three point to the fact that studying the features of labor market income is key to understand consumption and wealth inequality.

The distributions of consumption and wealth are of great importance to both academics and policymakers. While the former is closely related with the cross-sectional distribution of welfare ([Blundell and Preston \(1998\)](#)), the latter reflects more persistent differences across households, is more unequal, and has important implications for growth, mobility, and many other economic, social, and political considerations.

In this thesis, I have shown that household earnings face infrequent but relatively large shocks, that large negative shocks are more likely than large positive earnings shocks, and that persistence and earnings volatility vary significantly over the age and the earnings distribution. I have also shown that these features of earnings have changed over time, and that, particularly in the case of the Netherlands, the government substantially insures households against these fluctuations. Jointly with these facts on earnings, I have shown that also homeownership and the composition of household portfolios have changed over time, which has had an impact on the distribution of wealth within the bottom 95% of the wealth distribution.

I have then proposed two structural models that incorporate these important features to explain, respectively, the growth of consumption inequality over the life cycle and these changes in homeownership and household portfolios. The first of these models, in Chapter 2, is based on a very simple asset structure, but delivers a key message: households are aware that their earnings can suffer relatively large shocks and partially self-insure against these fluctuations. The second model, in Chapter 3, is much richer in its asset composition. It suggests that these changes in earnings dynamics over generations, together with changes in asset re-

turns and financial conditions, can rationalize the observed decrease in homeownership rates without assuming that younger generations do not want to buy houses anymore. Thus, it provides useful lessons for both academic research and policymakers who might be worried about the drop in homeownership rates.

Additionally, this thesis provides several technical contributions. It suggests, in Chapter 2, an easy way of incorporating rich household earnings dynamics into macroeconomic models. In Chapter 3, it enriches that process with business cycle and intergenerational variations. Furthermore, in Chapter 3 I propose and solve a life-cycle model which is richer in its asset and risk structure than it has been previously done in the literature.

I look forward to developing these themes further in my research agenda, and hope that the contributions in this thesis are useful for future researchers, policymakers, and the general public debate.

References

- Abowd, John and Card, David (1989), ‘On the covariance structure of earnings and hours changes’, *Econometrica* **57**(2), 411–45.
- Acemoglu, Daron and Autor, David (2011), Skills, tasks and technologies: Implications for employment and earnings, *in* ‘Handbook of labor economics’, Vol. 4, Elsevier, pp. 1043–1171.
- Aiyagari, S. R. (1994), ‘Uninsured idiosyncratic risk, and aggregate saving’, *Quarterly Journal of Economics* **109**(3), 659–684.
- Alan, Sule (2006), ‘Entry costs and stock market participation over the life cycle’, *Review of Economic Dynamics* **9**(4), 588–611.
- Altonji, Joseph G., Smith, Anthony A. and Vidangos, Ivan (2013), ‘Modeling earnings dynamics’, *Econometrica* **81**(4), 1395–1454.
- Ameriks, John and Zeldes, Stephen P (2004), ‘How do household portfolio shares vary with age’, Working paper.
- Arellano, Manuel, Blundell, Richard and Bonhomme, Stéphane (2017), ‘Earnings and consumption dynamics: a nonlinear panel data framework’, *Econometrica* **85**(3), 693–734.
- Attanasio, Orazio P. and Pavoni, Nicola (2011), ‘Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption’, *Econometrica* **79**(4), 1027–1068.
- Attanasio, Orazio P, Bottazzi, Renata, Low, Hamish W, Nesheim, Lars and Wakefield, Matthew (2012), ‘Modelling the demand for housing over the life cycle’, *Review of Economic Dynamics* **15**(1), 1–18.
- Bajari, Patrick, Chan, Phoebe, Krueger, Dirk and Miller, Daniel (2013), ‘A dynamic model of housing demand: Estimation and policy implications’, *International Economic Review* **54**(2), 409–442.
- Bansal, Ravi and Yaron, Amir (2004), ‘Risks for the long run: A potential resolution of asset pricing puzzles’, *The journal of Finance* **59**(4), 1481–1509.

- Barro, Robert J (2006), ‘Rare disasters and asset markets in the twentieth century’, *The Quarterly Journal of Economics* **121**(3), 823–866.
- Barro, Robert J (2009), ‘Rare disasters, asset prices, and welfare costs’, *American Economic Review* **99**(1), 243–64.
- Becker, Thomas A and Shabani, Reza (2010), ‘Outstanding debt and the household portfolio’, *Review of Financial Studies* **23**(7), 2900–2934.
- Bell, Felicitie C., Wade, Alice H. and Goss, Stephen C. (1992), ‘Life tables for the United States Social Security Area: 1900-2080’, (Social Security Administration, Office of the Actuary).
- Benabou, Roland (2002), ‘Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?’, *Econometrica* **70**(2), 481–517.
- Benzoni, Luca, Collin-Dufresne, Pierre and Goldstein, Robert S (2007), ‘Portfolio choice over the life-cycle when the stock and labor markets are cointegrated’, *The Journal of Finance* **62**(5), 2123–2167.
- Berger, David, Guerrieri, Veronica, Lorenzoni, Guido and Vavra, Joseph (2017), ‘House prices and consumer spending’, *The Review of Economic Studies* **85**(3), 1502–1542.
- Bewley, Truman (1977), ‘The permanent income hypothesis: A theoretical formulation’, *Journal of Economic Theory* **16**(2), 252–292.
- Blackburn, John O (1967), ‘Implicit tax rate reductions with growth, progressive taxes, constant progressivity, and a fixed public share’, *The American Economic Review* **57**(1), 162–169.
- Blundell, Richard and Preston, Ian (1998), ‘Consumption inequality and income uncertainty’, *The Quarterly Journal of Economics* **113**(2), 603–640.
- Blundell, Richard, Graber, Michael and Mogstad, Magne (2015), ‘Labor income dynamics and the insurance from taxes, transfers and the family’, *Journal of Public Economics* **127**, 58–73.

- Blundell, Richard, Pistaferri, Luigi and Preston, Ian (2008), ‘Consumption inequality and partial insurance’, *The American Economic Review* pp. 1887–1921.
- Blundell, Richard, Pistaferri, Luigi and Saporta-Eksten, Itay (2016), ‘Consumption inequality and family labor supply’, *The American Economic Review* **106**(2), 387–435.
- Bonaparte, Y, Korniotis, GM and Kumar, A (2018), ‘Portfolio choice and asset pricing with investor entry and exit’.
- Bonhomme, Stéphane and Robin, Jean-Marc (2009), ‘Assessing the equalizing force of mobility using short panels: France, 1990–2000’, *The Review of Economic Studies* **76**(1), 63–92.
- Borella, Margherita, De Nardi, Mariacristina and Yang, Fang (2019), The lost ones: the opportunities and outcomes of non-college educated americans born in the 1960s, Technical report, NBER Working Paper.
- Browning, Martin, Ejrnaes, Mette and Álvarez, Javier (2010), ‘Modelling income processes with lots of heterogeneity’, *Review of Economic Studies* **77**, 1353–1381.
- Busch, Christopher and Ludwig, Alexander (2017), Higher-order income risk over the business cycle: A parametric approach, Technical report, mimeo.
- Busch, Christopher, Domeij, David, Guvenen, Fatih and Madera, Rocio (2018), Asymmetric business-cycle risk and social insurance, Working Paper 24569, National Bureau of Economic Research.
- Cagetti, Marco and De Nardi, Mariacristina (2008), ‘Wealth inequality: Data and models’, *Macroeconomic Dynamics* **12**(S2), 285–313.
- Cagetti, Marco and De Nardi, Mariacristina (2009), ‘Estate taxation, entrepreneurship, and wealth’, *American Economic Review* **99**(1), 85–111.
- Campanale, Claudio, Fugazza, Carolina and Gomes, Francisco (2015), ‘Life-cycle portfolio choice with liquid and illiquid financial assets’, *Journal of Monetary Economics* **71**, 67–83.

- Campbell, John and Deaton, Angus (1989), ‘Why is consumption so smooth?’, *The Review of Economic Studies* **56**(3), 357–373.
URL: <http://dx.doi.org/10.2307/2297552>
- Campbell, John Y and Yogo, Motohiro (2006), ‘Efficient tests of stock return predictability’, *Journal of financial economics* **81**(1), 27–60.
- Castañeda, Ana, Díaz-Giménez, Javier and Ríos-Rull, José-Víctor (2003), ‘Accounting for the U.S. Earnings and Wealth Inequality’, *Journal of Political Economy* **111**(4), pp. 818–857.
- Chambers, Matthew, Garriga, Carlos and Schlagenhaut, Don E (2009), ‘Accounting for changes in the homeownership rate’, *International Economic Review* **50**(3), 677–726.
- Chang, Minsu (2018), ‘A house without a ring: The role of changing marital transitions for housing decisions’, Working paper.
- Chetty, Raj, Sándor, László and Szeidl, Adam (2017), ‘The effect of housing on portfolio choice’, *The Journal of Finance* **72**(3), 1171–1212.
- Civale, Simone, Díez-Catalán, Luis and Fazilet, Fatih (2016), ‘Discretizing a process with non-zero skewness and high kurtosis’, Mimeo, University of Minnesota.
- Cocco, Joao F (2005), ‘Portfolio choice in the presence of housing’, *Review of Financial studies* **18**(2), 535–567.
- Cocco, Joao F, Gomes, Francisco J and Maenhout, Pascal J (2005), ‘Consumption and portfolio choice over the life cycle’, *Review of financial Studies* **18**(2), 491–533.
- Conesa, Juan Carlos, Kitao, Sagiri and Krueger, Dirk (2009), ‘Taxing capital? not a bad idea after all!’, *American Economic Review* **99**, 25–48.
- Davidoff, Thomas (2006), ‘Labor income, housing prices, and homeownership’, *Journal of urban Economics* **59**(2), 209–235.
- Davis, Morris A and Ortalo-Magné, François (2011), ‘Household expenditures, wages, rents’, *Review of Economic Dynamics* **14**(2), 248–261.

- Davis, Steven J and Von Wachter, Till M (2011), ‘Recessions and the cost of job loss’.
- De Nardi, Mariacristina (2004), ‘Wealth inequality and intergenerational links’, *Review of Economic Studies* **71**, 743–768.
- De Nardi, Mariacristina and Fella, Giulio (2017), ‘Saving and wealth inequality’, *Review of Economic Dynamics* **26**, 280 – 300.
- De Nardi, Mariacristina, Fella, Giulio and Paz-Pardo, Gonzalo (2016), ‘The implications of richer earnings dynamics for consumption and wealth’, NBER Working paper no. 21917. National Bureau of Economic Research.
- De Nardi, Mariacristina, Fella, Giulio and Paz-Pardo, Gonzalo (2019), ‘Non-linear household earnings dynamics, self-insurance, and welfare’, *Journal of the European Economic Association* . Forthcoming.
- De Nardi, Mariacristina, Fella, Giulio, Knoef, Marike, Paz-Pardo, Gonzalo and Van Ooijen, Ran (2018), ‘Family and government insurance: Wage, earnings, and income risks in the netherlands and the u.s.’, NBER working paper No. 25832.
- De Nardi, Mariacristina, French, Eric and Benson, David (2011), ‘Consumption and the great recession’.
- Deaton, Angus and Paxson, Christina (1994), ‘Intertemporal choice and inequality’, *Journal of Political Economy* **102**(3), 437–67.
- Deelen, Anja and Euwals, Rob (2014), ‘Do wages continue increasing at older ages? evidence on the wage cushion in the netherlands’, *De Economist* **162**(4), 433–460.
- Díaz, Antonia and Luengo-Prado, María José (2008), ‘On the user cost and homeownership’, *Review of Economic Dynamics* **11**(3), 584–613.
- De Nardi, Mariacristina, French, Eric and Jones, John B. (2010), ‘Why do the elderly save? The role of medical expenses’, *Journal of Political Economy* **118**(1), 39–75.

- Duca, John V, Muellbauer, John and Murphy, Anthony (2011), ‘House prices and credit constraints: Making sense of the us experience’, *The Economic Journal* **121**(552), 533–551.
- Dynan, Karen E, Elmendorf, Douglas W and Sichel, Daniel E (2006), ‘Can financial innovation help to explain the reduced volatility of economic activity?’, *Journal of monetary Economics* **53**(1), 123–150.
- Epstein, Larry G and Zin, Stanley E (1989), ‘Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework’, *Econometrica: Journal of the Econometric Society* pp. 937–969.
- Fagereng, Andreas, Gottlieb, Charles and Guiso, Luigi (2017), ‘Asset market participation and portfolio choice over the life-cycle’, *The Journal of Finance* **72**(2), 705–750.
- Fagereng, Andreas, Guiso, Luigi and Pistaferri, Luigi (2016), ‘Portfolio choices, firm shocks and uninsurable wage risk’, Working paper.
- Fagereng, Andreas, Holm, Martin Blomhoff, Moll, Benjamin and Natvik, Gisle (2019), ‘Saving behavior across the wealth distribution: The importance of capital gains’, Working paper.
- Falcattoni, Elena and Schmitz, James (2018), ‘Public policies to alleviate the housing crisis’, mimeo.
- Fella, Giulio (2014), ‘A generalized endogenous grid method for non-smooth and non-concave problems’, *Review of Economic Dynamics* **17**, 329–344.
- Fella, Giulio, Gallipoli, Giovanni and Pan, Jutong (2017), ‘Markov-chain approximations for life-cycle models’.
- Fischer, Marcel and Khorunzhina, Natalia (2019), ‘Housing decision with divorce risk’, *International Economic Review* .
- Fisher, Jonas DM and Gervais, Martin (2011), ‘Why has home ownership fallen among the young?’, *International Economic Review* **52**(3), 883–912.

- Flavin, Marjorie and Yamashita, Takashi (2011), ‘Owner-occupied housing: life-cycle implications for the household portfolio’, *The American Economic Review* **101**(3), 609–614.
- Gálvez, Julio (2017), Household portfolio choices and nonlinear income risk, Technical report, Working paper.
- Gerardi, Kristopher, Rosen, Harvey S and Willen, Paul (2007), Do households benefit from financial deregulation and innovation? the case of the mortgage market, Technical report, National Bureau of Economic Research.
- Gervais, Martin (2002), ‘Housing taxation and capital accumulation’, *Journal of Monetary Economics* **49**(7), 1461–1489.
- Geweke, John and Keane, Michael (2000), ‘An empirical analysis of earnings dynamics among men in the psid: 1968-1989’, *Journal of Econometrics* **96**, 293–356.
- Glaeser, Edward L, Gottlieb, Joshua D and Gyourko, Joseph (2012), Can cheap credit explain the housing boom?, *in* ‘Housing and the financial crisis’, University of Chicago Press, pp. 301–359.
- Goldin, Claudia Dale and Katz, Lawrence F (2009), *The race between education and technology*, Harvard University Press.
- Golosov, Mikhail, Troshkin, Maxim and Tsyvinski, Aleh (2016), ‘Redistribution and social insurance’, *American Economic Review* **106**(2), 359–86.
- Gomes, Francisco and Michaelides, Alexander (2005), ‘Optimal life-cycle asset allocation: Understanding the empirical evidence’, *The Journal of Finance* **60**(2), 869–904.
- Graber, Michael and Lise, Jeremy (2015), ‘Labor market frictions, human capital accumulation and consumption inequality’, mimeo, University College London.
- Guvenen, Fatih (2007), ‘Learning your earning: Are labor income shocks really very persistent?’, *The American Economic Review* **97**(3), 3.

- Guvenen, Fatih and Smith, Anthony A (2014), ‘Inferring labor income risk and partial insurance from economic choices’, *Econometrica* **82**(6), 2085–2129.
- Guvenen, Fatih, Kaplan, Greg, Song, Jae and Weidner, Justin (2017), ‘Lifetime incomes in the united states over six decades’.
- Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar and Song, Jae (2016), What do data on millions of U.S. workers reveal about life-cycle earnings risk? Working paper, University of Minnesota.
- Guvenen, Fatih, Ozkan, Serdar and Song, Jae (2014), ‘The nature of countercyclical income risk’, *Journal of Political Economy* **122**(3), pp. 621–660.
- Haider, S. J. (2001), ‘Earnings instability and earnings inequality of males in the united states: 1967-1991’, *Journal of Labor Economics* **19**(4), 799–836.
- Halvorsen, Elin, Holter, Hans, Ozkan, Srdar and Storesletten, Kjetil (2019), Dissecting idyosincertaic income risk. Mimeo.
- Heathcote, Jonathan, Perri, Fabrizio and Violante, Giovanni L (2010), ‘Unequal we stand: An empirical analysis of economic inequality in the united states, 1967-2006’, *Review of Economic dynamics* **13**(1), 15–51.
- Heathcote, Jonathan, Storesletten, Kjetil and Violante, Giovanni L (2005), ‘Two views of inequality over the life cycle’, *Journal of the European Economic Association* **3**(2-3), 765–775.
- Heathcote, Jonathan, Storesletten, Kjetil and Violante, Giovanni L. (2010), ‘The macroeconomic implications of rising wage inequality in the united states’, *Journal of Political Economy* **118**(4), 681–722.
- Heathcote, Jonathan, Storesletten, Kjetil and Violante, Giovanni L (2014), ‘Optimal tax progressivity: An analytical framework’, National Bureau of Economic Research.

- Hoffman, Eran B. and Malacrino, Davide (2019), ‘Employment time and the cyclical growth of earnings’, *Journal of Public Economics* **169**, 160–171.
- Huggett, Mark, Ventura, Gustavo and Yaron, Amir (2011), ‘Sources of lifetime inequality’, *American Economic Review* **101**(7), 2923–2954.
- Hurd, Michael D and Smith, James P (2001), Anticipated and actual bequests, in ‘Themes in the Economics of Aging’, University of Chicago Press, pp. 357–392.
- Jacobson, Louis S., LaLonde, Robert J. and Sullivan, Daniel G. (1993), ‘Earnings losses of displaced workers’, *The American Economic Review* **83**(4), 685–709.
- Jäntti, Markus and Jenkins, Stephen P. (2015), Income mobility, in Anthony B. Atkinson and François Bourguignon, eds, ‘Handbook of Income Distribution’, Vol. 2 of *Handbook of Income Distribution*, Elsevier, pp. 807–935.
- Jappelli, Tullio and Pistaferri, Luigi (2006), ‘Intertemporal choice and consumption mobility’, *Journal of the European Economic Association* **4**(1), 75–115.
- Johnson, Kathleen W and Li, Geng (2010), ‘The debt-payment-to-income ratio as an indicator of borrowing constraints: Evidence from two household surveys’, *Journal of Money, Credit and Banking* **42**(7), 1373–1390.
- Kaplan, Greg (2012), ‘Inequality and the life cycle’, *Quantitative Economics* **3**(3), 471–525.
- Kaplan, Greg and Violante, Giovanni L. (2010), ‘How much consumption insurance beyond self-insurance?’, *American Economic Journal: Macroeconomics* **2**(4), 53–87.
- Kaplan, Greg and Violante, Giovanni L (2014), ‘A model of the consumption response to fiscal stimulus payments’, *Econometrica* **82**(4), 1199–1239.

- Karahan, Faith and Ozkan, Serdar (2013), ‘On the persistence of income shocks over the life cycle: Evidence, theory and implications’, *Review of Economic Dynamics* **16**(3), 452–476.
- Kennan, John (2006), ‘A note on discrete approximations of continuous distributions’, *University of Wisconsin Madison* pp. 211–251.
- Krueger, Dirk and Perri, Fabrizio (2003), ‘On the welfare consequences of the increase in inequality in the united states’, *NBER macroeconomics annual* **18**, 83–121.
- Krueger, Dirk and Perri, Fabrizio (2006), ‘Does income inequality lead to consumption inequality? evidence and theory’, *The Review of Economic Studies* **73**(1), 163–193.
- Krueger, Dirk, Mitman, Kurt and Perri, Fabrizio (2016), ‘Macroeconomics and household heterogeneity’, *Handbook of Macroeconomics* **2**, 843–921.
- Kuhn, Moritz and Ríos-Rull, José-Víctor (2015), ‘2013 update on the us earnings, income and wealth distributional facts: a view from macroeconomics’, Working paper.
- Kuhn, Moritz, Schularick, Moritz and Steins, Ulrike I (2017), ‘Income and wealth inequality in america, 1949-2013’, Working paper.
- Li, Wenli, Liu, Haiyong, Yang, Fang and Yao, Rui (2015), ‘Housing over time and over the life cycle: A structural estimation’, FRB of Philadelphia Working Paper.
- Lillard, Lee A. and Willis, Robert J. (1978), ‘Dynamic aspects of earning mobility’, *Econometrica* **46**(5), 985–1012.
- Lippi, Francesco and Perri, Fabrizio (2019), ‘Unequal growth’, mimeo.
- Lovenheim, Michael F (2011), ‘The effect of liquid housing wealth on college enrollment’, *Journal of Labor Economics* **29**(4), 741–771.
- Lundberg, Shelly and Pollak, Robert A (2007), ‘The american family and family economics’, *Journal of Economic Perspectives* **21**(2), 3–26.

- McKay, Alisdair (2017), ‘Time-varying idiosyncratic risk and aggregate consumption dynamics’, *Journal of Monetary Economics* **88**, 1 – 14.
- McLean, R David and Pontiff, Jeffrey (2016), ‘Does academic research destroy stock return predictability?’, *The Journal of Finance* **71**(1), 5–32.
- Meghir, Costas and Pistaferri, Luigi (2004), ‘Income variance dynamics and heterogeneity’, *Econometrica* **72**(1), 1–32.
- Meghir, Costas and Pistaferri, Luigi (2011), ‘Earnings, consumption and life cycle choices’, *Handbook of Labor Economics* **4**, 773–854.
- Mitchell, Olivia S and Phillips, John WR (2006), ‘Social security replacement rates for alternative earnings benchmarks’.
- Nakajima, Makoto (2005), ‘Rising earnings instability, portfolio choice, and housing prices’, *Unpublished Manuscript* .
- Nakajima, Makoto (2010), ‘Optimal capital income taxation with housing’, Working paper.
- Paz-Pardo, Gonzalo (2019), ‘Homeownership and portfolio choice over the generations’.
- Pesaran, M Hashem and Timmermann, Allan (1995), ‘Predictability of stock returns: Robustness and economic significance’, *The Journal of Finance* **50**(4), 1201–1228.
- Primiceri, Giorgio E and Van Rens, Thijs (2009), ‘Heterogeneous life-cycle profiles, income risk and consumption inequality’, *Journal of Monetary Economics* **56**(1), 20–39.
- Pruitt, S. and Turner, N. (2018), The nature of household labor income risk. Finance and Economics Discussion series 2018-034.
- Quadrini, Vincenzo and Ríos-Rull, José-Víctor (2014), Inequality in macroeconomics, in A.B. Atkinson and F. Bourguignon, eds, ‘Handbook of Income Distribution, vol. 2B’, Elsevier, chapter 15.

- Ravn, Morten O and Sterk, Vincent (2017), ‘Job uncertainty and deep recessions’, *Journal of Monetary Economics* **90**, 125–141.
- Rouwenhorst, Geert (1995), Asset pricing implications of equilibrium business cycle models, *in* Thomas F.Cooley, ed., ‘Frontiers of business cycle research’, Princeton University Press, chapter 10.
- Saiz, Albert (2010), ‘The geographic determinants of housing supply’, *The Quarterly Journal of Economics* **125**(3), 1253–1296.
- Scholz, John Karl, Seshadri, Ananth and Khitatrakun, Surachai (2006), ‘Are americans saving “optimally” for retirement?’, *Journal of political economy* **114**(4), 607–643.
- Shiller, Robert J (2015), *Irrational exuberance: Revised and expanded third edition*, Princeton University Press.
- Smith, Lawrence B, Rosen, Kenneth T and Fallis, George (1988), ‘Recent developments in economic models of housing markets’, *Journal of economic literature* **26**(1), 29–64.
- Storesletten, Kjetil, Telmer, Chris I. and Yaron, Amir (2004a), ‘Cyclical dynamics in idiosyncratic labor-market risk’, *Journal of Political Economy* **112**, 695–717.
- Storesletten, Kjetil, Telmer, Christopher I. and Yaron, Amir (2004b), ‘Consumption and risk sharing over the life cycle’, *Journal of Monetary Economics* **51**(3), 609–633.
- Tauchen, George (1986), ‘Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions’, *Economics Letters* **20**(2), 177–181.
- Vestman, Roine (2012), ‘Limited stock market participation among renters and home owners’.
- Vissing-Jorgensen, Annette (2002), ‘Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures’, Technical report, National Bureau of Economic Research.

Yao, Rui and Zhang, Harold H (2005), 'Optimal consumption and portfolio choices with risky housing and borrowing constraints', *Review of Financial studies* **18**(1), 197–239.

A Appendix to Chapter 1

1.1 Dutch male wages, computed using actual hours worked.

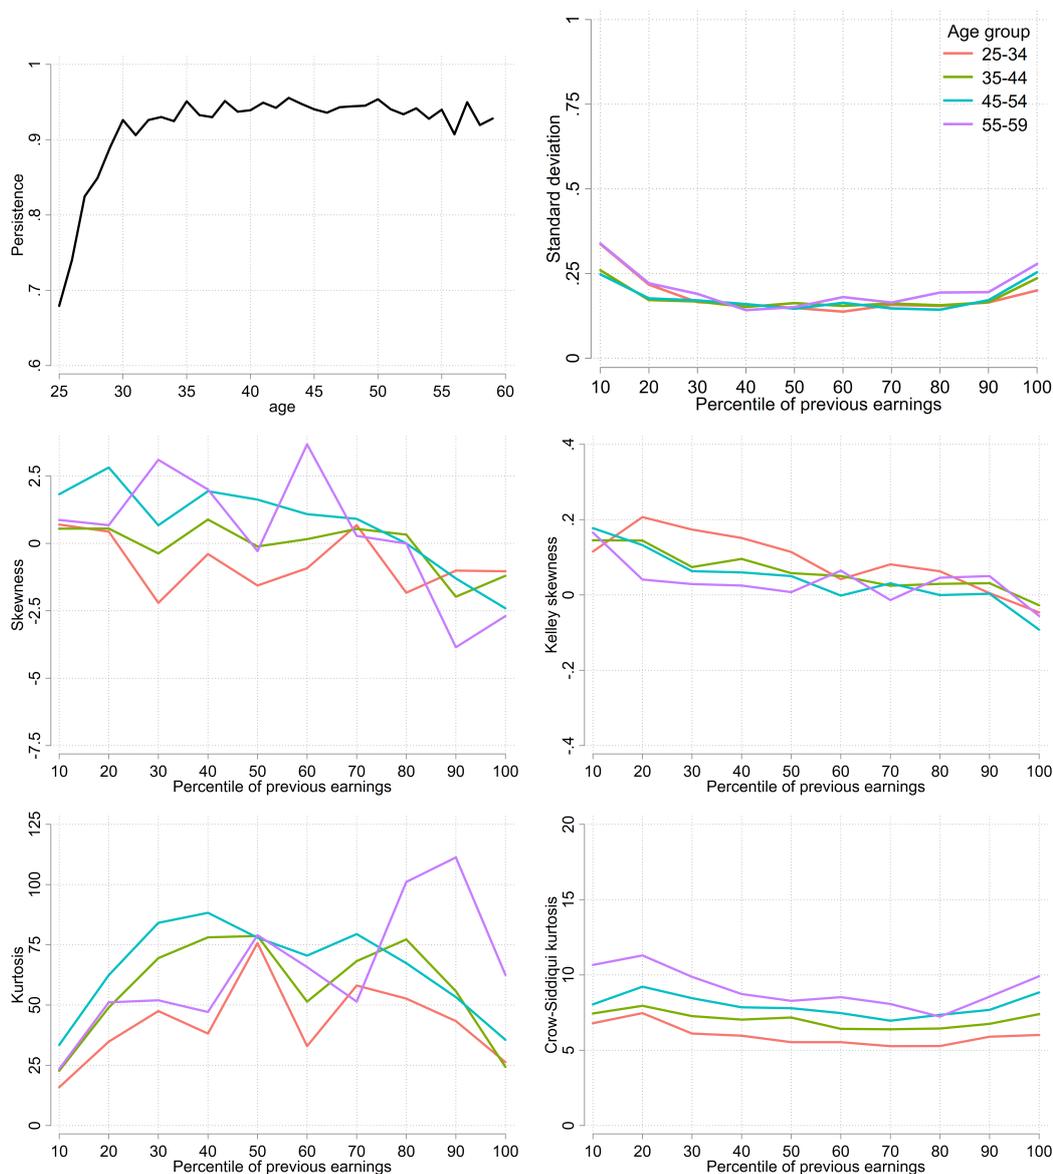


Figure A.1: Dutch male wages, computed using actual hours worked. Wage persistence (top left) and following moments of wage changes: standard deviation (top right), skewness (middle left), Kelley's skewness (middle right), kurtosis (bottom left), and Crow-Siddiqui kurtosis (middle right), by age group and previous earnings percentile.

1.2 Non-robust measures of skewness and kurtosis

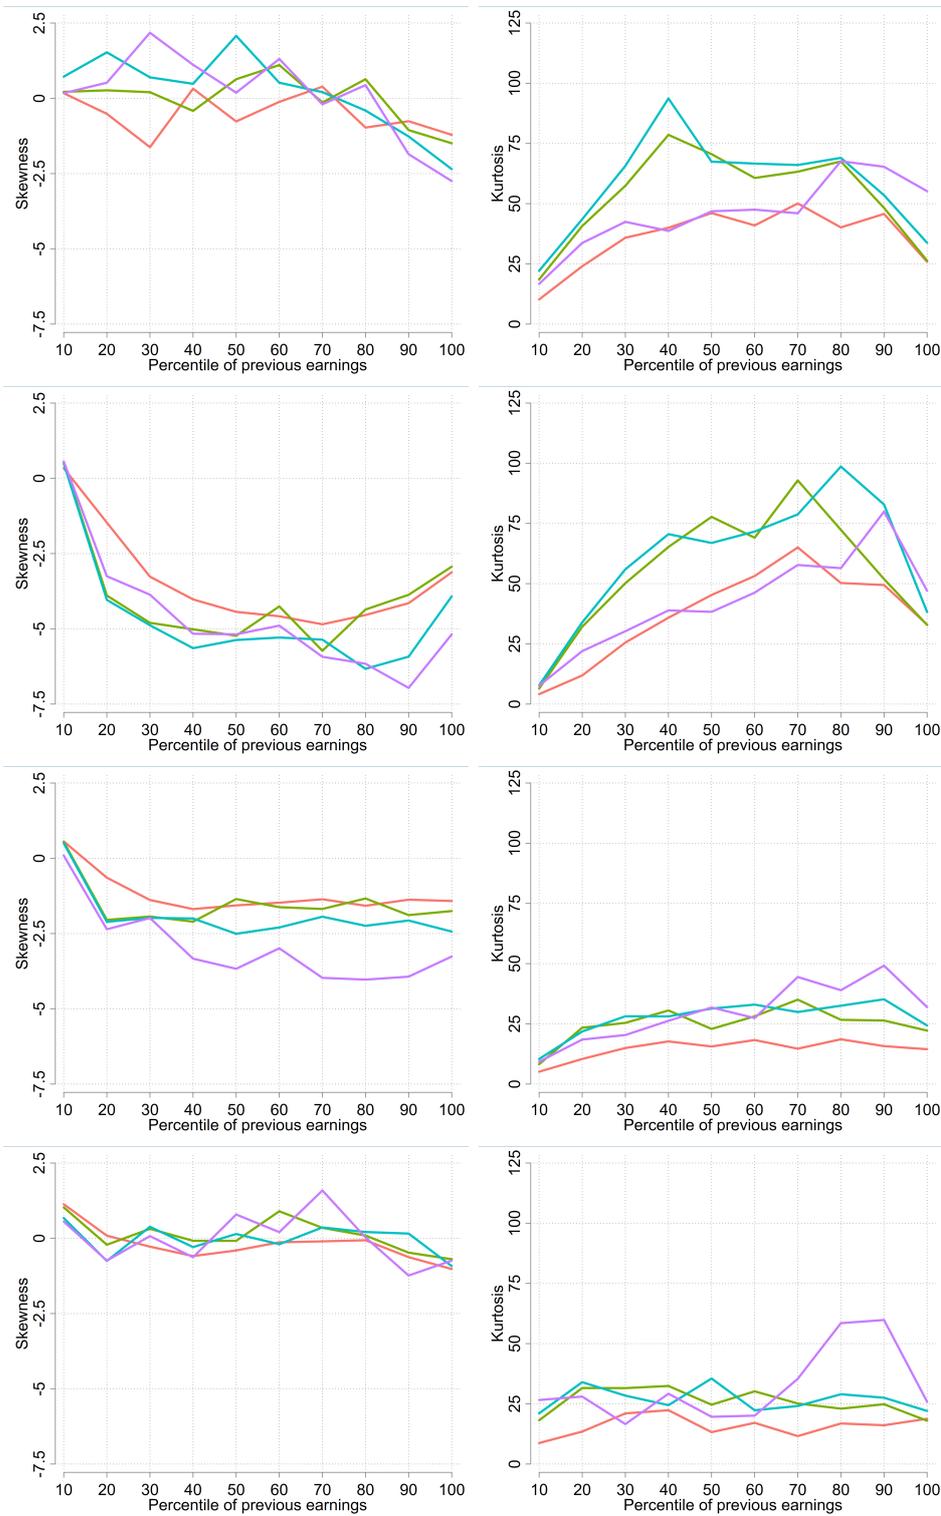


Figure A.2: The Dutch data: Non-robust measures of Skewness (left) and Kurtosis (right): male wages (first row) male earnings (second row) before-tax household income (third row) after-tax household income (fourth row).

1.3 Male earnings by recent earnings

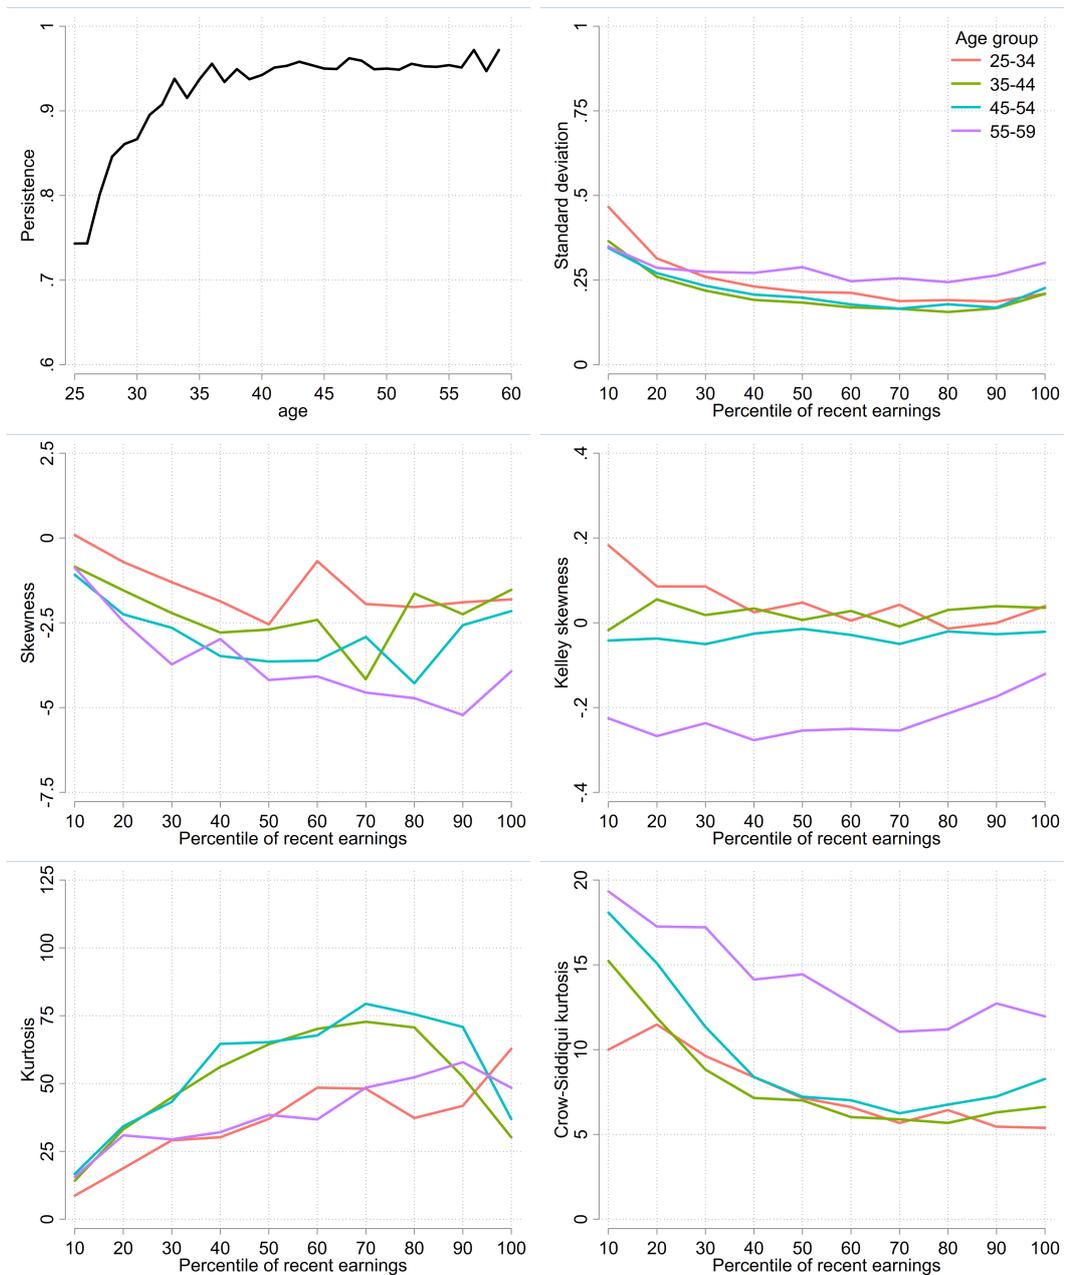
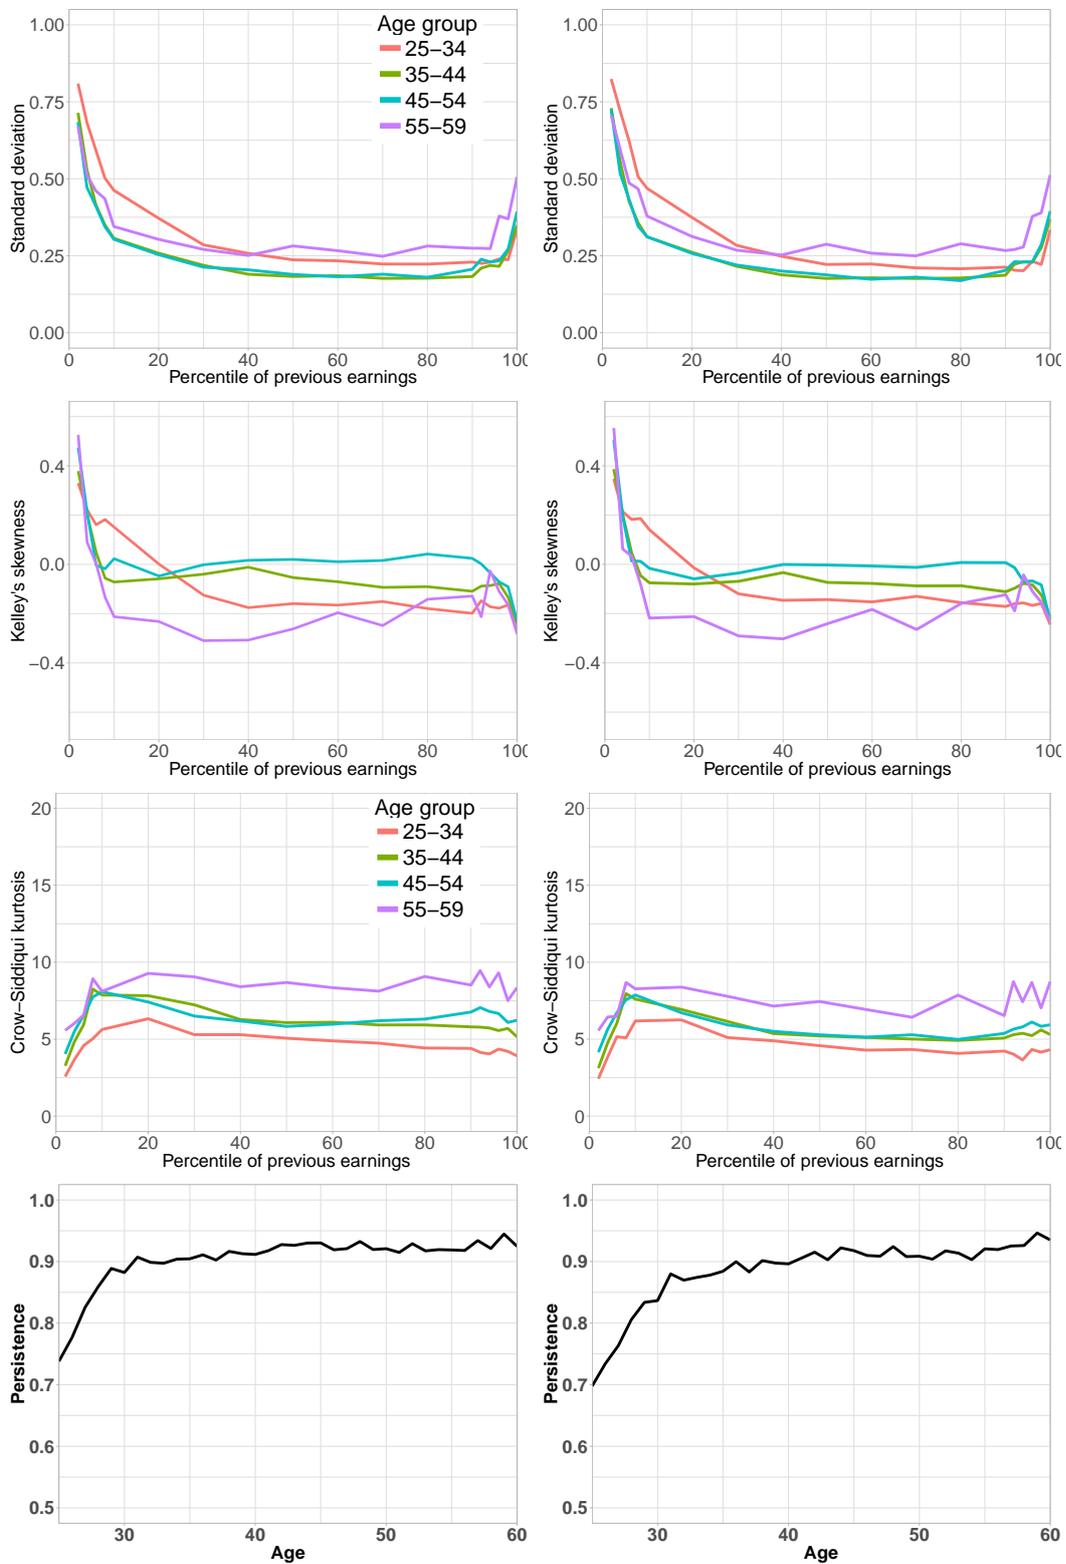


Figure A.3: The Dutch data: Male earnings by recent earnings.

1.4 Household earnings and pre-tax income

Figure A.4 compares household earnings with household pre-tax income, which is the sum of household earnings and capital income. The figures are almost the same, indicating that allowing for capital income makes little difference.

Figure A.4: The Dutch data: household earnings (left) and household pre-tax income (right)



B Appendix to Chapter 2

2.1 PSID data

We use PSID data to estimate both the canonical and the nonlinear earnings processes. In this Appendix we briefly describe the PSID, our sample selection criterion, the precise variable definition we use and the details of the estimation of the canonical earnings process.

2.1.1 The PSID

The Panel Study of Income Dynamics (PSID) follows a large number of U.S. households over time and reports information about their demographic characteristics and sources of income. The PSID was initially composed of two major subsamples. The first of them, the SRC (Survey Research Center) or core subsample, was designed to be representative of the U.S. population and is a random sample itself, including over 18,000 individuals living in 5,000 households. The second, the SEO (Survey of Economic Opportunity) subsample, was created to study the characteristics of the most deprived households. Later, Immigrant and Latino subsamples were also added to the PSID.

From 1968 to 1997, the survey was yearly. After 1997, it started having a biennial structure. We only consider the SRC or core subsample because the SEO oversamples the poor. After dropping the SEO and Latino samples we are left with a random sample, which makes computations simpler since weights are not needed (Haider, 2001).⁷⁰

2.1.2 Sample Selection

Since the model period is one year, we restrict ourselves to the yearly part of the survey, and focus on the years 1968-1992. We have dropped the 1993-1997 period because there was a major redesign of the survey in 1993. This affected the method through which information was collected, with the introduction of computer-based surveys, and the definitions of

⁷⁰It must be taken into account that the weighting of our final dataset can be affected by attrition and by the fact that we are neglecting observations of yearly income under \$ 1500 (expressed in 2015 dollars).

some variables we use (for instance, asset income of other family members was no longer available, and wife labor income was redefined). We have verified that results are not sensitive to including these five years.

Following standard practice in the literature, we only consider individuals between ages 25 and 60. We consider all households, whether or not male-headed. This differs from many other papers, but follows e.g. [Krueger, Mitman and Perri \(2016\)](#).

We deflate values to 2013 dollars, and only keep observations with earnings (as defined below) above \$1500. This is also in accordance with standard practice in the literature, where observations below a minimum earnings threshold are dropped ([De Nardi \(2004\)](#) or [Guvenen et al. \(2016\)](#), for instance).

2.1.3 Income Definition

For the estimation of our earnings process, we use after-tax equivalized household earnings. We first construct nonfinancial pre-tax household earnings using PSID data. Then, we estimate a tax function to obtain after-tax earnings. Finally, we regress earnings on the number of family members for the purposes of equivalization.

We now describe each of these three steps in detail.

2.1.3.1 Nonfinancial Pre-Tax Income

We construct nonfinancial pre-tax income closely following [Guvenen and Smith \(2014\)](#). The procedure is based on subtracting all asset income from total family income.

Before 1976, asset income was not directly available in the survey. Therefore, we take total family income, subtract head and wife taxable income (which includes labor and asset income) and then add back labor earnings for head and wife independently.

From 1976, we consider all the available measures of asset income. These include farm income, business income, rent and interests, with the addition of gardening and roomers income (from 1978), and asset income of family members other than head and wife (from 1984).

We keep top-coded observations, but drop the very small number (8) of households who, probably due to measurement error, would have non-financial income below zero.

2.1.3.2 Tax Function

We obtain **disposable** labor income by subtracting an estimated measure of taxes on labor income from nonfinancial pre-tax income.

We first compute the total amount of income taxes paid by households by adding up the federal income tax variable (which is available in the PSID until 1990) with a constructed measure of payroll (FICA and Medicare) taxes, which is based on applying the historical rates and caps to labor earnings of husband and wife independently.

We then separate taxes on labor and asset income by running a regression of this total tax measure on nonfinancial income and its square, and asset income and its square. This also follows [Güvenen and Smith \(2014\)](#). The estimated coefficients allow us to predict taxes on labor income⁷¹, which we subtract from nonfinancial pre-tax income to get after-tax labor income.

2.1.3.3 Equivalization

We then equivalize after-tax nonfinancial disposable income by running a regression of earnings on the number of family members and keeping the residuals. We also extract year fixed effects.

2.1.4 Estimating the Canonical and Nonlinear Earnings Processes

2.1.4.1 Estimating the Nonlinear Earnings Process

To finally implement the [Arellano et al. \(2017\)](#) procedure, we create a sample with all sets of subsequent three-year observations (without replacement: once an observation in the PSID sample is in a 3-year set in our sample we drop it). This implies that we are also dropping all of those

⁷¹We use the coefficients estimated in the sample 1968-1990 to predict taxes on labor income for the period 1968-1992, given that the PSID federal income tax variable is not available for the last two years of our sample.

households that do not have three consecutive valid income observations in the PSID.

We then follow the procedure described in Section 2.4 and the discretization explained in Section 2.5.

2.1.4.2 Estimating the Canonical Earnings Process

In Storesletten et al. (2004b) (and in many other papers in the literature, e.g. Krueger et al. (2016)) the earnings process is estimated by fitting a parametric process to the variance of earnings profile that we observe in the data. The standard way is to compute the variance of earnings by age-cohort-year cells, and then get the coefficients of a regression of those on either age and year or age and cohort. For consistence with our approach and with the consumption data we rely on, we use the one that controls for year effects (see discussion below).

We follow a GMM procedure in which we minimize the distance of the estimated process to the profile of variances and first-order autocovariances of earnings over the life cycle⁷². The weighting matrix is the identity matrix.

The canonical earnings process in equations (6)-(8) implies (for $t > 1$)

$$y_{it} = \rho^{t-1}\eta_{i1} + \sum_{j=2}^t \rho^{t-j}\zeta_{ij} + \epsilon_{it} \quad (\text{B.1})$$

from which

$$\text{var}(y_{it}) = \rho^{2(t-1)}\sigma_{\eta_1}^2 + \sum_{j=2}^t \rho^{2(t-j)}\sigma_{\zeta}^2 + \sigma_{\epsilon}^2. \quad (\text{B.2})$$

and

$$\text{cov}(y_{it}, y_{i,t+1}) = \rho^{2t-1}\sigma_{\eta_1}^2 + \sum_{j=2}^t \rho^{1+2(t-j)}\sigma_{\zeta}^2 \quad (\text{B.3})$$

follow, allowing to identify moments.

2.2 Computation of the Variances of Log Earnings and Log Consumption.

We estimate the canonical earnings process described in Section 2.4.1 by matching the variance and first autocovariance of log earnings.

⁷²We describe in Appendix 2.2 how we compute these variances.

To compute the variance of log earnings, which we report in Figure 13, we use the procedure described in Kaplan (2012) (Appendix C.3), controlling for year effects, with our PSID data. More specifically, we take log disposable and equivalized labor income \tilde{y}_{it} , where i indexes the household and t is the age of its head, and run the regression

$$\tilde{y}_{it} = \beta_t' \mathfrak{D}_t + \beta_d' \mathfrak{D}_d + y_{it}, \quad (\text{B.4})$$

where \mathfrak{D}_t and \mathfrak{D}_d are matrices with columns corresponding to a full set of age and year (date) dummies, respectively. The vectors β_t and β_d are the corresponding coefficients and y_{it} the earnings residuals.⁷³

We compute the variance of y_{it} by age group as

$$\text{Var}_t(y) = \frac{1}{D} \sum_{d=1}^D \left(\sum_{i=1}^{N_{d,t}} \frac{y_{it}^2}{N_{d,t}} \right), \quad (\text{B.5})$$

where D is the number of years in the dataset, and $N_{d,t}$ is the numerosity of each age-year cell. This implies that the variance of earnings at age t weighs equally the corresponding conditional variances of earnings in each year.

We also compute the variance of y_{it} by age group controlling for cohort instead of year effect, using the cohort counterpart of equation (B.5)

$$\text{Var}_t(y) = \frac{1}{K(t)} \sum_{k=1}^{K(t)} \left(\sum_{i=1}^{N_{k,t}} \frac{y_{it}^2}{N_{k,t}} \right), \quad (\text{B.6})$$

where $K(t)$ is the number of cohorts containing individuals of age t and $N_{k,t}$ is the numerosity of each cohort-age cell.⁷⁴ This approach weighs the conditional variances from each cohort equally.

Under both approaches, we obtain very similar age profiles (Figure B.1) and parameter estimates for the canonical process (Table B.1).

Turning to consumption, we compute the variance of log consumption using data from the CEX for the period 1980-2007. Nondurable consump-

⁷³As described in Appendix B, we use the earnings residuals from equations (B.4) to estimate our earnings processes.

⁷⁴The residuals used in equations (B.5) and (B.6) are the same. Given that year, age and cohort are linearly dependent, the residuals from equation (B.4) are the same that would obtain from projecting onto age and cohort dummies.

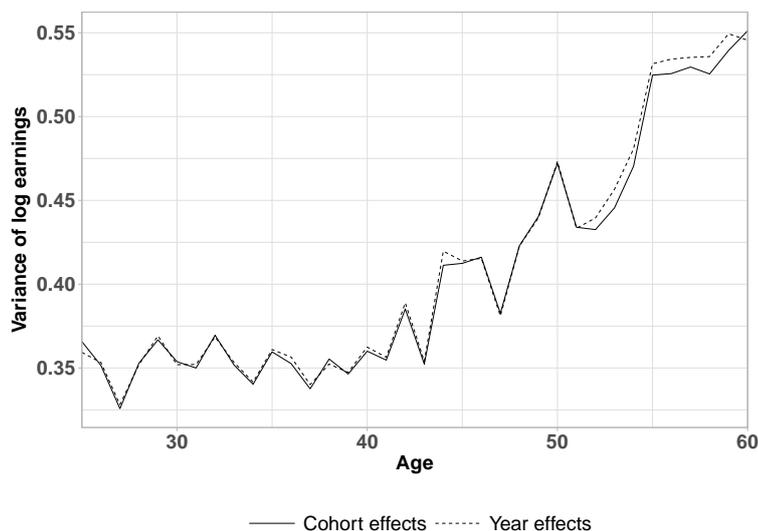


Figure B.1: Cross-sectional variance of log earnings over the life cycle, cohort effects vs year effects

	σ_{ϵ}^2	$\sigma_{\eta_1}^2$	σ_{ζ}^2	ρ
Year effects	0.0620	0.2332	0.0060	1
Cohort effects	0.0669	0.2379	0.0057	1

Table B.1: Estimates for the canonical earnings process: cohort vs. year effects

tion includes food, clothing, gasoline, household operation, transportation, medical care, recreation, tobacco, and education.

We compute the variance of log consumption following the same procedures that we use for the variance of log earnings. Namely, we deal with year effects using the method proposed by [Kaplan \(2012\)](#) and equalize consumption with a regression on the number of family members. We have also applied this procedure to OECD-equivalized consumption data and verified that it yields very similar results to [Heathcote, Perri and Violante \(2010\)](#) when they control for year effects.

2.3 Robust Measures of Skewness and Kurtosis

Figure [B.2](#) represents a robust measure of skewness (Kelley’s skewness) and a robust measure of kurtosis (Crow-Siddiqui kurtosis) for male pre-tax earnings in the W2, and for both male pre-tax and household after-tax earnings in the PSID.

Equations B.7 and B.8 show that these measures are computed taking into account specific percentiles of the distribution of earnings changes and, as such, are robust to the effect of outliers. The inspection of B.2 shows that all of the main features highlighted in Section 2.2 are still present in these more robust measures.

$$KS = \frac{(P_{90} - P_{50}) - (P_{50} - P_{10})}{P_{90} - P_{10}} \quad (\text{B.7})$$

$$CS = \frac{P_{97.5} - P_{2.5}}{P_{75} - P_{25}} \quad (\text{B.8})$$

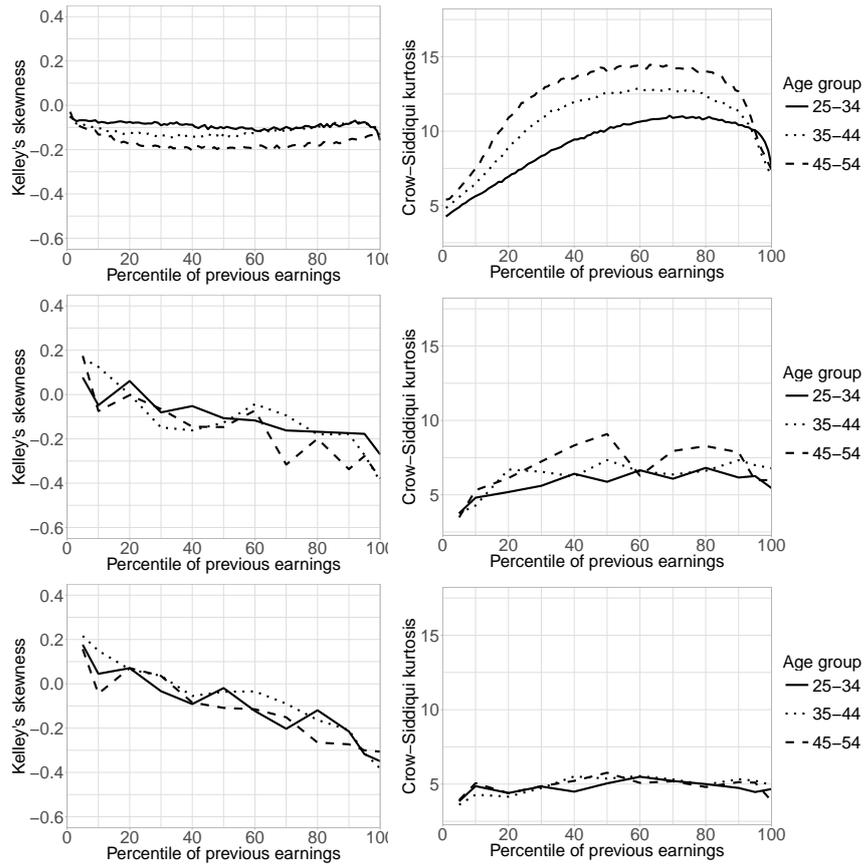


Figure B.2: Kelly skewness and Crow-Siddiqui kurtosis of male pre-tax earnings growth in the W2 (top panel) and PSID (central panel), and of household after-tax earnings growth in the PSID (bottom panel)

C Appendix to Chapter 3

3.1 Data

The data used in this paper is taken mostly from the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF). The former is particularly valuable because it allows to follow several cohorts of households over a very long period of time (1968-2013), and it is the base for the estimation of the earnings process and most of the housing-related measures. However, it contains, particularly in its earlier periods, very limited information about stock market participation, wealth, and financial asset allocation of households. For these measures I rely on the Survey of Consumer Finances, which contains very detailed information about households' balance sheets. However, with limited exceptions, it lacks a panel dimension.

I now briefly describe the characteristics of each of the surveys, the sample selection criteria, and the estimation of the several targets and profiles used in the paper.

3.1.1 PSID

The Panel Study of Income Dynamics (PSID) follows a large number of U.S. households and their qualifying spinoffs since 1968 and provides information about demographic characteristics, sources of income, housing status, and, since more recently, their wealth and consumption. When it started, the PSID was composed of two main subsamples: the SRC (Survey Research Center), which was designed to be representative of the U.S. population at the time and which is a random sample itself, and the SEO (Survey of Economic Opportunity), which oversampled the poor. Later, the PSID was augmented with the Immigrant and Latino subsamples.

The survey was yearly from 1968 to 1997, and started being biennial since then. Wealth information was available in the 1984, 1989, 1994 waves, and from 1999 onwards, and has become progressively richer and improved in quality. Since 1999 it broadly replicates the wealth inequality patterns present in the SCF without oversampling the richest.

For my main results, I drop the SEO, Latino and Immigrant samples and am therefore left with a random sample, which makes computations simpler given that weights are not needed (Haider, 2001). However, the weighting of the final dataset can be affected by attrition and the sample selection requirements. In Appendix 3.3 I report how different sample selection and deflation procedures affect some key features of the earnings process.

3.1.1.1 Measures and sample selection

In the PSID data, I define cohorts as follows:

- The **1940s cohort** are the households whose head was born between 1940 and 1950. For the estimation of the income process, I increase the sample size and consider households between 1930 and 1950. For simulation purposes, I consider they were born in 1942.
- The **1960s cohort** are households whose head was born between 1960 and 1970 (1950-1970 for the income process). For simulation purposes, I consider they were born in 1962.
- The **1980s cohort** are households whose head was born between 1980 and 1990 (1970-1990 for the income process). For simulation purposes, I consider they were born in 1982.

Naturally, the changes in earnings dynamics, homeownership, etc. have happened progressively over time and do not necessarily correspond with the admittedly arbitrary boundaries set. To some extent, the features of the 1940s earnings process are a weighted average of the earnings process of people born between 1930 and 1950. However, including those allows me to increase the sample sizes, and to obtain more observations of people going through a recession at different ages.

For the earnings process, I assume that all business cycle effects are absorbed by the business-cycle dependent process, so I do not extract year effects. For the representation of changes in earnings risk and the computation of age-efficiency profiles, I extract a linear yearly trend from earnings data.

The estimation of the household income process requires eliminating households that display very low attachment to the labor market (whose labor income in a given year is below a minimum level of \$1500 in 2013 prices). This assumption is standard in the earnings processes literature (see [De Nardi et al. \(2019\)](#) for details) and avoids issues related with taking logs of very small numbers. Furthermore, I also drop those households for whom there are no two consecutive observations available.

For the older ages of the younger cohorts, there are some cases in which there are very few observations for a particular combination of the labor market aggregate state, which are not sufficient for the estimation of the flexible parameters. In those cases I replace the missing cohort-age-states with their correspondent levels of the previous cohort. This affects the 1960s cohort after age 50 for the states of recovery and staying in a recession, and the 1980s cohort for all states related to a recession after age 30. I follow a similar procedure for years which are not yet observable (1980s cohort from ages 35-40), and provide some robustness checks with respect to this assumption in [Section 3.7.2](#).

For all of the other measures reported in the paper (homeownership, etc.) none of these restrictions are imposed. In particular, I do not require the sample to be composed of the same households in every year. This allows me to keep a bigger sample, but implies the assumption that any attrition or nonresponse happens randomly and does not affect the evolution of the measures reported.

With respect to the two types of housing, “detached houses” are those defined in the PSID as “detached single family houses” and “non-detached houses” are all other types of structures (including 2-family houses, apartments, etc.).

From 1968 to the end of the sample, housing PTI ratios are computed as the ratio of the median house price reported by PSID homeowners to median household income in the PSID. In order to estimate the evolution of housing PTI ratios before the start of the PSID (as agents are born with age 20 in the model, they live in the model from 1962 to 1967), I use the housing data provided by Robert Shiller ([Shiller, 2015](#)) and assume the PTI to be above trend whenever the house price index is above its

trend and viceversa.

3.1.1.2 Tax progressivity

As explained in Section 3.4.4, the model explicitly includes the choice between taking the standard deduction or itemized deductions. However, this implies that the tax progressivity coefficient τ in Equation 43 needs to be reestimated, because in previous studies, such as [Heathcote et al. \(2014\)](#), it was computed taking into account the existence of itemization and the standard deduction. Removing the standard deduction from disposable income implies a reduction of the progressivity coefficient, as it is an important driver of progressivity in the US and other tax codes ([Blackburn, 1967](#)).

To perform this estimation, I need to compute the counterfactual disposable income or, alternatively, the counterfactual level of taxes paid T_{it}^* by a household with pre-tax income y_{it} in the absence of standard deduction or HMID. To do so, I first estimate the following equation in my PSID sample:

$$\log \hat{y}_{it} = \lambda_1 + (1 - \tau_1) \log y_{it} \quad (\text{C.1})$$

where \hat{y}_{it} is post-tax pre-benefit household income. I use the estimated parameters from this equation to predict T_{it}^* , assuming that $T_{it}^* = \log \tilde{y}_{it}^* - \lambda_1 + (1 - \tau_1) \log y_{it}^*$, where $y_{it}^* = y_{it} + \max(sd, HM)$, sd stands for the standard deduction, and HM for the HMID that corresponds to a given household. The basic assumption here is that the taxes paid by a household with a certain level of income in the counterfactual world with no standard deduction are the same as those paid by a household with that level of income plus the deduction in the observed world. As for the mortgage deductions, I define them to be the product of the average mortgage interest rate in a certain year and the outstanding mortgage the household claims to have in the PSID, as long as they are smaller than the total mortgage payments the household has made in the previous year.

Once I have counterfactual taxes T_{it}^* , I construct counterfactual disposable income (post-tax, post-benefit) $\tilde{y}_{it}^* = \tilde{y}_{it} + T_{it} - T_{it}^*$, run the following

regression:

$$\log \tilde{y}_{it}^* = \lambda + (1 - \tau) \log y_{it} \quad (\text{C.2})$$

and obtain $\tau = 0.085$.

3.1.2 SCF

The Survey of Consumer Finances (SCF), conducted every three years since 1983, provides information about the financial situation of US households. It contains detailed data on household balance sheets, income, and other demographic characteristics. Given the focus on wealth, the survey oversamples the rich, who hold most of the assets in the economy. To do so, it combines an area-probability sample (geographical stratification) with a list sample that guarantees that a sufficient amount of wealthy individuals are included.

Apart from the 1983-2007 waves, I also consider the older historical waves of the Survey of Consumer Finances made available by the University of Michigan. Namely, I consider the 1963, 1968, 1969, 1970, 1971, and 1977 waves to construct, where relevant, statistics like wealth to income ratios or shares of stock market participants. While this data is less exhaustive than the recent waves of the SCF, it is the only source to provide reliable information about household wealth and its composition for the cohorts I am interested in before 1980. For a longer discussion, analysis and harmonization procedures of these waves of the SCF, see [Kuhn, Schularick and Steins \(2017\)](#).

I use the SCF to obtain stock market participation data. I consider a household to be participating in the stock market if any of its members hold stocks directly or indirectly via an IRA account, Keogh plan, mutual fund or pension plan like a 401(k). For most of the waves the composition of these funds is provided as a categorical variable, so for indirect holders I consider them to be stockholders if they hold any such plan which is formed of “mostly or all stock”.

3.1.2.1 Bequest targets

The main bequest targets are based on [Hurd and Smith \(2001\)](#), who use the Asset and Health Dynamics among the Oldest Old (AHEAD) study from the Health and Retirement Study (HRS). This study focuses on households whose heads were born in 1923 or before, which is significantly earlier than the first cohort considered in this paper. Therefore, the bequest targets reflect bequests left by a generation which had potentially different characteristics to the ones I consider.

Bequest data for the 1940s cohort is not yet available. However, SCF data reveals that the generation born in the 1940s had accumulated around 30% more wealth at age 75 than the generation born around 1910-1920. Thus, for the main version of this model I adjust the average bequest target by increasing it by exactly as much as average wealth increased at age 75 between these cohorts. This imperfectly captures several possible reasons for holding more wealth during retirement (longer life expectancies, different patterns of medical expenditure, different histories) - future data on wealth decumulation by this cohort can impose more discipline on this assumption. I do not adjust the targeted percentage of people with zero bequests, but the calibration tends to underpredict it.

3.1.3 Higher order moments of earnings

I define skewness and kurtosis as the third and fourth standardized moments of log earnings changes, respectively. Kelley's skewness (KS) and Crow-Siddiqui kurtosis (CS) are defined following [Güvenen et al. \(2016\)](#):

$$KS = \frac{(P_{90} - P_{50}) - (P_{50} - P_{10})}{P_{90} - P_{10}} \quad (C.3)$$

$$CS = \frac{P_{97.5} - P_{2.5}}{P_{75} - P_{25}} \quad (C.4)$$

where P represents a percentile in the distribution of earnings changes. Thus, Kelley's skewness is more positive the further away the 90th percentile is from the median; and more negative the further away the 10th percentile is from the median, while Crow-Siddiqui kurtosis is larger the fatter the tails of the distribution are. I refer to both as *robust* mea-

asures because they are less affected by outliers than standard skewness and kurtosis.

3.2 Computational appendix

3.2.1 EGM method

For a given set of parameters, the model is solved using a combination of value function iteration (VFI) and the endogenous gridpoint method (EGM), following the algorithm described in [Fella \(2014\)](#).

The model has several discrete states (housing, aggregate state) and continuous states (safe assets, risky assets, earnings). After the last period of life (age $T + 1$), utility is known as it can be directly derived from the bequest function (30). From then I proceed via backwards induction.

For any age t , given a set of states (y, a, f, h, m, Ω) we need to find the policy functions for the household for current period's consumption c and next period's safe assets a' , risky assets f' , housing h' , and mortgages m' . Using standard methods, this would imply computing the value associated with each of the feasible choices and maximizing over the 4-D space (as we can always solve for one of the choices using the budget constraint) to find the optimal choice for each set of states.

Using the EGM method allows to substantially speed up the computation of a pair of these choices. For this paper, I use the EGM method to solve the (c, a') choice conditional on the f', h' choice. For computational purposes, $m = -a$, and I allow for the grid of f to include some points that correspond to possible positive holdings of a . Given f' and h' , I use the inverted Euler equation to compute the consumption choice c that corresponds to each future choice of assets a' .

By the budget constraint, the sum of consumption and all savings must equal current cash on hand. This means that, given f', h' and the pair c, a' we have found with the Euler equation, we can interpolate the endogenous grid of consumption to the exogenous grid of cash on hand that is determined by the states in period t , and obtain the f', h' -conditional choices of c and a' for those particular states.

The only point left is to then run a nonlinear maximizer over the f'

and h' choices, thus obtaining all required policy functions. Using this procedure, the nonlinear maximizer is only run over a 2-D space, which implies the algorithm is significantly faster.

The Euler equation does not necessarily hold in all scenarios for safe assets (for instance, when the household is borrowing constrained with no housing, or at the boundary of the condition that requires it to pay interest on its debt). These situations are dealt with specifically.

3.2.2 Global minimization

I solve the model using FORTRAN 2008. Due to the large state space, it is very computationally intensive - in a workstation with 44 cores it takes roughly 25 minutes to solve for a given parametrization. In order to find the parameter values that minimize the weighted square distance between the targets and their values in the data, I use a modified version of the NEWUOA numerical optimization algorithm.

I acknowledge the use of the UCL Myriad High Throughput Computing Facility (Myriad@UCL), and associated support services, in the completion of this work.

3.3 Intergenerational differences, alternative definitions

The changes across cohorts described in Section 3.2 are based on the SRC, which is the PSID's representative sample of the 1968 US population and their offspring, are deflated with the CPI, and consider household earnings for all households, whether married or not. In this Section, I describe the qualitative and quantitative implications of considering three alternative approaches: picking the whole PSID and weighting it (Appendix 3.3.1), deflating with the PCE (Appendix 3.3.2), and selecting married households only (Appendix 3.3.3).

3.3.1 Sample composition

In this section, I make use of the whole PSID rather than the SRC. Because only the SRC is a random sample, this also requires making use of the

PSID-provided weights. In the earlier years of the survey, this means considering the Survey of Economic Opportunity part of the PSID, that oversamples the poor; in later years of the survey, in particular after 1997, it implies taking into account the immigrant population that has arrived to the US since. Using the SRC rather than the whole PSID for the main results leads to a cleaner comparison - keeping the offspring of the same population of reference helps to ascribe the changes in labor market dynamics to structural changes in the labor market, opportunities, and family formation, rather than changing demographics across the whole society. However, it has the disadvantage on missing out on additional population growth of potentially different socioeconomic characteristics, which has implications for house prices and more broadly any general equilibrium effects.

Figure C.1 shows median and average earnings for this broader sample. Including immigrants reduces the median earnings of younger cohorts with respect to the oldest generation.

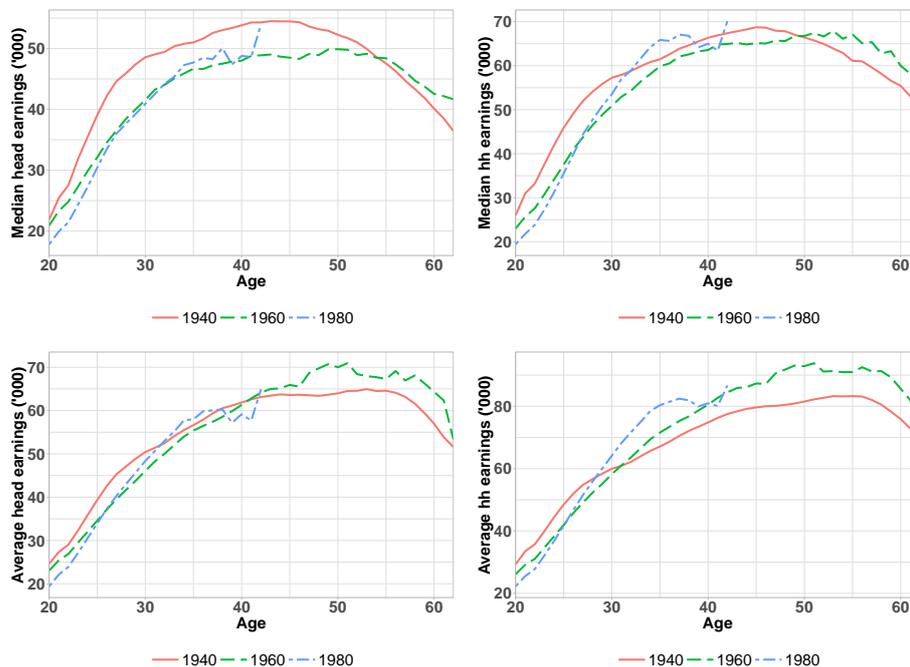


Figure C.1: Median earnings (top) and average earnings (bottom), for household heads (left) and households (right), 2013 dollars, sample including SEO and immigrants.

Turning to distributional features, we observe in Figure C.2 that considering all immigrants has the interesting implication that the difference

between the 1940s and 1960s cohort is preserved, or even larger than before, but the difference between the 1960s and the 1980s cohort becomes almost insignificant. While earnings have become more unequal for the sons and daughters of the original PSID sample members, the entrance of immigrants has contributed to reduce the variance of the earnings distribution of the 1980s cohort with respect to the 1960s.

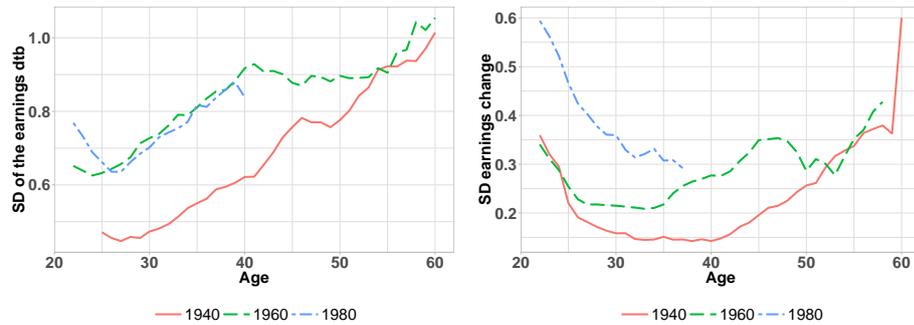


Figure C.2: Standard deviation of log earnings (left) and earnings changes (right), sample including SEO and immigrants

The measure of homeownership considered in this paper can also depend on sample choices, particularly if immigrants and newer incorporations to the PSID sample have substantially different homeownership patterns. Figure C.3 shows the resulting comparison. All main patterns are similar: if anything, taking into account the whole population implies a marginally larger gap between the 1940s and 1960s cohort, and a marginally smaller gap between the 1960s and 1980s cohort, which is consistent with the smaller gap in earnings inequality and risk in this wider sample.

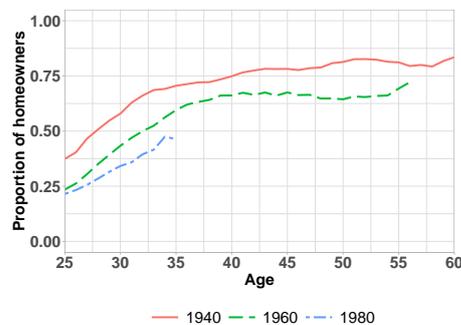


Figure C.3: Homeownership, by cohort, sample including SEO and immigrants

3.3.2 Choice of deflator

In this section, I deflate earnings with the PCE or personal consumption expenditure deflator rather than the CPI. These two measures differ slightly on their scope and their computation procedure. While the PCE takes into account all expenditure made by households and also on behalf of households, such as total medical expenditures, the CPI only considers what households spend out-of-pocket. The PCE is based on business surveys, while the reference basket for the CPI is based on data from the Consumer Expenditure Survey or CEX. Given that the focus of the paper refers to the consumption and portfolio possibilities of all but the richest of households, the main results are deflated with the CPI, which more closely reflects the changes in prices of the goods and services that households actually pay. The PCE index is instead more frequently used when performing aggregate macroeconomic analysis.

In this sample period, the PCE implies overall lower cumulative inflation than the CPI and, therefore, implies that median and average earnings of younger cohorts have grown more than with the CPI. However, because cross-sectional inequality within a cohort-age cell is not affected by the choice of deflator, the facts regarding changes in the distribution of earnings that lie at the core of this paper are unchanged when considering the PCE.

Figure C.4 shows median and average earnings by cohort, for male and head and spouse earnings. Deflating with the PCE increases the differences between cohorts, particularly for household earnings, which implies that it acts in the opposite direction as the inclusion of a broader sample described in Appendix 3.3.1. Naturally, the choice of deflator does not affect earnings inequality within age and cohort, nor measures of earnings risk, nor homeownership.

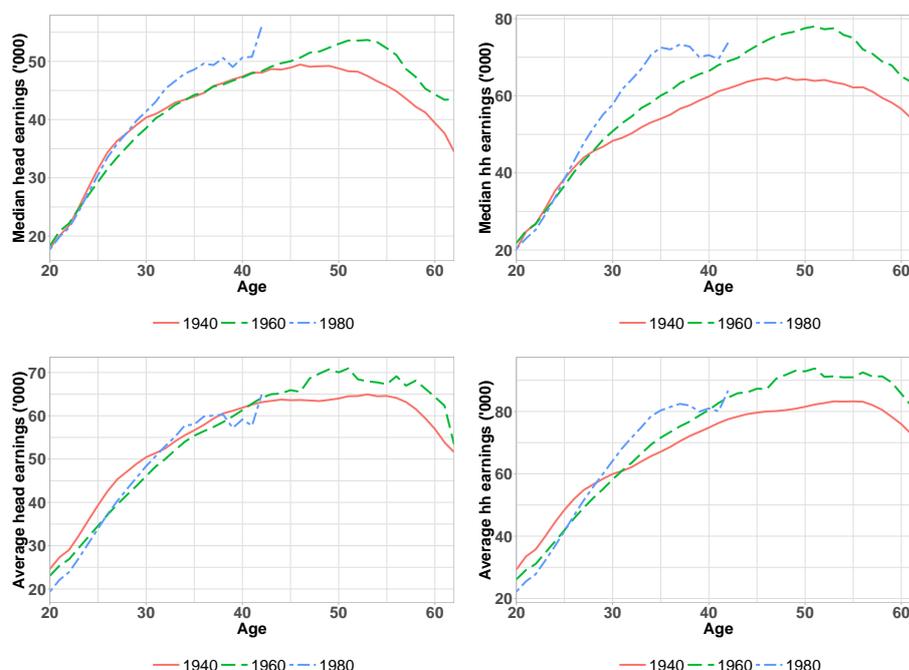


Figure C.4: Median (top) and average (bottom) earnings, 2013 dollars. Left: household heads, right: households. PCE-deflated sample.

3.3.3 Marital dynamics and family composition

Figure C.5 compares the profiles in Figure 22 with those for married households. While it is clear that family composition affects the profiles for the first ten years of age, and that earnings inequality is lower in the more homogeneous sample of married couples, the main picture is pretty similar, which suggests that there are differences in labor earnings across cohorts, particularly in distributional terms, which are not fully explained by the differential timing of marriage.

Figure C.6 shows that there are also large differences in homeownership over different cohorts if we restrict the sample to married households or to households with children, which provides additional evidence to suggest that earnings dynamics are relevant over and above changes in family composition. Naturally, in these selected samples homeownership tends to be larger than in the general population.

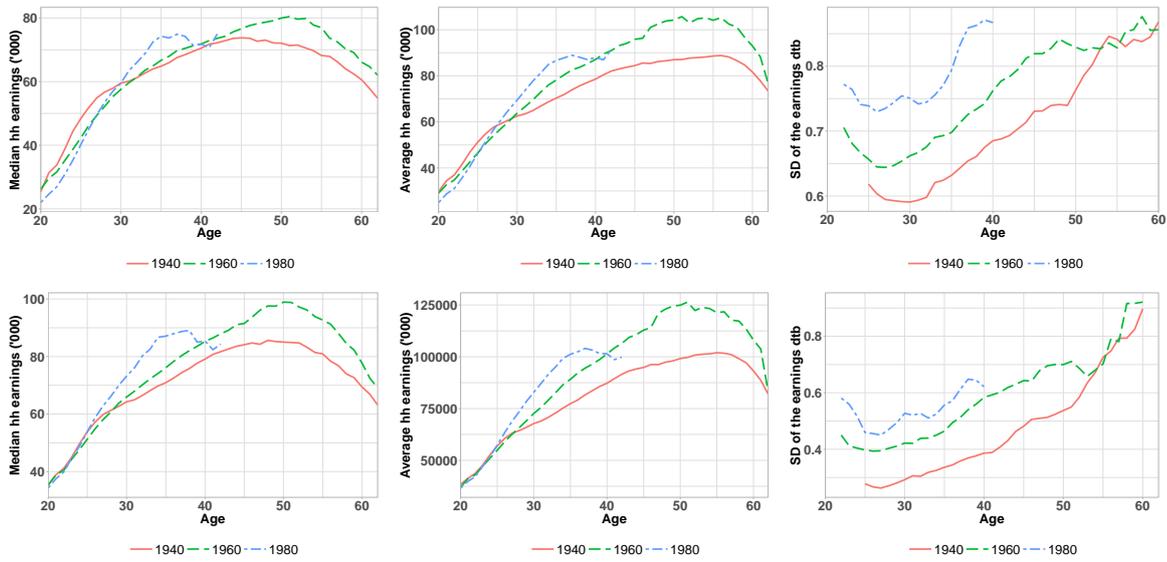


Figure C.5: Changes in the earnings distribution over the generations. Top: household earnings; bottom: household earnings for married households. Left: median earnings; center, average earnings; right: standard deviation of the log earnings distribution.

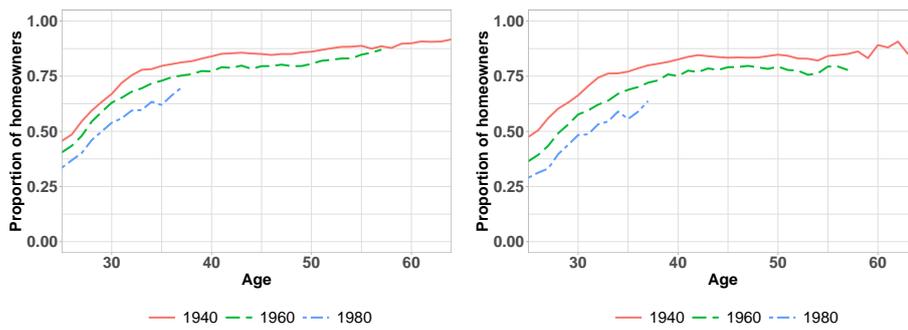


Figure C.6: Homeownership by cohorts, PSID data. Left, sample restricted to married couples; right, sample restricted to households with at least one child.

3.3.4 Earnings process, fit of variances

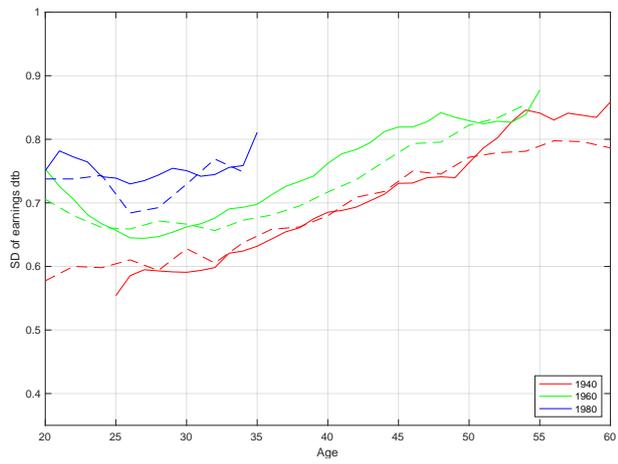


Figure C.7: Variance of log earnings over the life cycle, PSID data vs model-implied

3.3.5 Earnings process, additional Figures

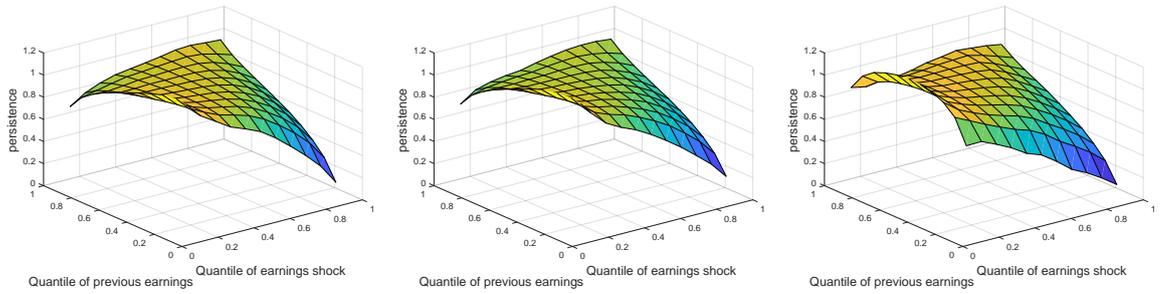


Figure C.8: Nonlinear persistence, by cohort. Top left: 1940s; top right: 1960s; bottom: 1980s. All ages

3.3.6 Portfolio composition, additional Figures

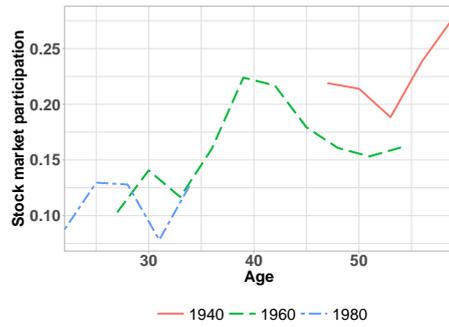


Figure C.9: Direct stock market participation, three generations. SCF data.

3.3.7 Equivalence scales by generation

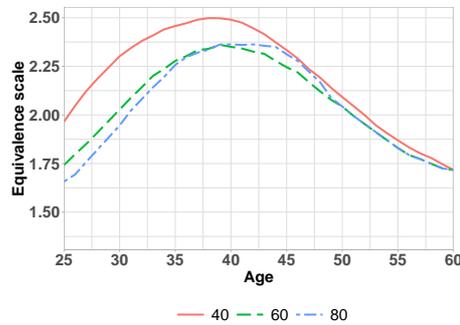


Figure C.10: Average OECD equivalence scale, by cohort

3.4 Model, additional results and descriptions

3.4.1 Adjusting housing prices

I compute the results in which housing prices are left free to adjust as a response to policy changes or earnings counterfactuals in the following way:

First, I begin by defining the *baseline stock of housing* H_t^s as the total number of houses that are occupied by their owners, which is equivalent to the number of households that live in owner-occupied housing. In order to aggregate across cohorts, I simulate a total of 31 birth-year cohorts (born in all even years from 1930 to 1990), and I attribute the earnings process of the 1940s cohort to the group 1930-1949, the earnings process of the 1960s cohort to the group 1950-1969, and the earnings process of the 1980s cohort to the group 1970-1990. Like in the main experiment,

each birth cohort experiments aggregate shocks as they happened in the data in each calendar year, that corresponds to a specific age for each of the 32 cohorts.

Given that I do not have information on the earnings processes of the relevant years for the cohorts born before 1930s, I make the simplifying assumption that their homeownership profiles are not impacted by any of the changes introduced in the paper and they hold a constant amount of housing. The impact of this assumption would depend on the reaction of these non-modelled cohorts, but it is only particularly restricting for the earlier cohorts when they are relatively young.

Thus, H_t^s is computed as summarized in Equation C.5, where I am aggregating over all households i that belong to each of these 31 birth-year cohorts. I then assume that, at given prices p_t^h , H_t^s is the total amount of housing supplied for these cohorts; which implies that p_t^h is the price that clears the market given housing supply and demand.

$$\int_i H_i^d(p_t^h) = H_t^s \tag{C.5}$$

Then, I am interested in computing how the quantities of housing bought by each of the different cohorts change as external factors, such as the earnings process, change. A change in the earnings process will induce a change in the housing demand functions H_i^{d*} for each household. To which extent this gets translated into changes in quantities exchanged and changes in prices depends on the elasticity of housing supply.

The main results in the paper are computed under the assumption that housing supply is fully elastic at given prices, and that prices would not respond to changes in the earnings process. I first compute the exact opposite case (all adjustments happening via prices), and then show an intermediate case in which the elasticity of housing supply is nonzero but finite. For all three cases, I begin by computing housing demand functions for each household, dependent on housing prices, at the new, counterfactual earnings process H^{d*} .

Assuming a fully inelastic housing supply implies fixing H_t^s at every year t and finding p_t^{h*} such that Equation C.6 holds. I do so sequentially

for $t = 1, \dots, Y$, where Y is the total amount of years in the simulation, and to that extent the model continues to capture history dependence of past prices. Furthermore, the slow-moving nature of these changes and the definition of H_t^s , which captures cyclical variations which the model is attributing to housing supply or other non-modelled factors, implies that the housing growth aggregate state is still consistent with the evolution of house prices.

$$\int_i H_i^{d*}(p_t^{h*}) = H_t^s \quad (\text{C.6})$$

Assuming an elastic housing supply implies, at a given (empirical) housing supply elasticity η , that the price that clears the market $p_t^{h'}$ satisfies:

$$\int_i H_i^{d*}(p_t^{h'}) = H_t^{s'} \quad (\text{C.7})$$

such that

$$\frac{\frac{p_t^{h'} - p_t^h}{p_t^h}}{\frac{H_t^{s'} - H_t^s}{H_t^s}} = \eta \quad (\text{C.8})$$

By substituting C.8 into C.7, one can solve for $p_t^{h'}$ in every period.

These experiments represent an approximation to actual equilibrium determination of housing prices. On the one hand, I do not model the agents that make housing supply decisions and approximate them with an isoelastic function (with elasticity which is either zero or 1.75). On the other hand, from the perspective of households, housing prices are still exogenous shocks with the same process as in the main version of the model. They are not aware that house prices are determined in equilibrium in a way that depends on total housing demand, and they do not perceive that intergenerational changes could be affecting house prices.

Figure C.11 shows the results for the case in which housing supply elasticity is assumed to be zero.

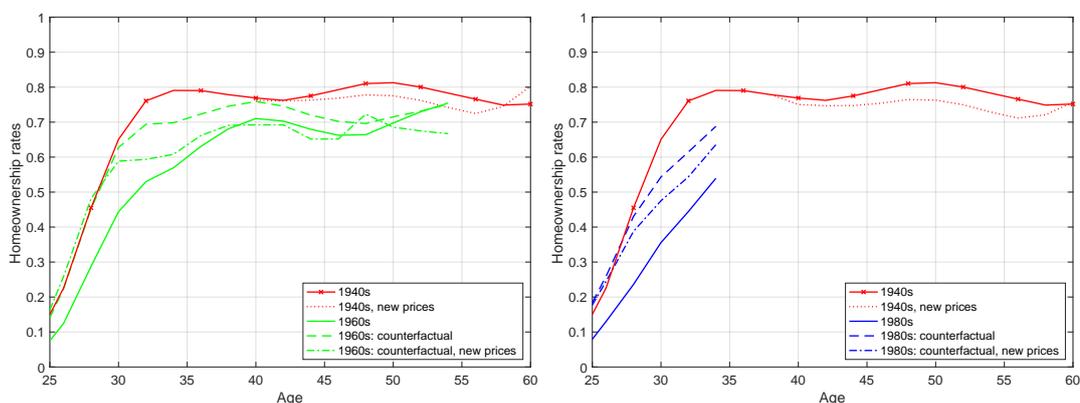


Figure C.11: Homeownership by cohorts, benchmark vs. counterfactual earnings processes. Assuming fully inelastic housing supply.

3.4.2 Household reactions to aggregate shocks

The realizations of the aggregate state directly affect household wealth and portfolio composition, and also household expectations about the future, which impact household decisions on their consumption of nondurable goods and housing services.

In order to understand the reactions of households to realizations of the aggregate state in the context of the Great Recession (Section 3.7.1.2), I simulate cohorts in a similar manner to that described in Section 3.4.1 and aggregate over them for each calendar year. Given that I am using 20-year windows for this simulation, this implies that I can exactly reproduce the histories for everyone who was 18 to 68 years old in 2008 and, as argued in Sections 3.5.3, 3.6.1, and 3.6.2, the model does well in generating their portfolio positions.

I then perform counterfactual experiments with respect to the baseline case in which 2008 was identical to 2006: high house prices, expectations about house price growth, good stock returns, and no recession, and then I turn on and off each element at a time. In the case in which I include the observed drop in stock returns in the analysis, the effects on homeownership are similar to the main case but the drop in consumption is substantially larger. This large drop is related to the assumption that all households have the same returns on stocks, to the lack of international diversification, and to the biennial nature of the model, which implies that households can only adjust their portfolio holdings every two years.

3.4.3 Canonical process

Figure C.12 shows the fit of the life-cycle variances for the canonical process with business cycle variation described in Section 3.8.1, and Table C.1 shows its estimated parameters.

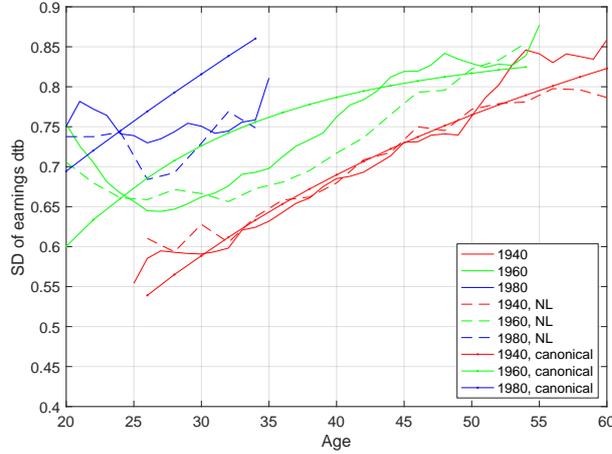


Figure C.12: Variance of earnings over the life cycle. Solid lines: data; dashed lines: NL process; dash-dot lines: canonical process

Cohort	ρ	$\sigma_{r,c}^2$	$\sigma_{b,c}^2$	σ_{ϵ}^2
1940	1.0	0.0155	0.0118	0.32
1960	0.99	0.0010	0.0198	0.35
1980	0.45	0.5796	0.6259	0.0001

Table C.1: Parameter estimates, business-cycle varying canonical process

The estimates for the 1940s cohort are more precisely estimated and thus as expected (variances are larger in recessions than in expansions). For the 1960s cohort, the difference between expansions and recessions is quite imprecisely estimated, but the model successfully replicates the higher level of earnings inequality that the cohort faces. Finally, for the 1980s cohort there are very few years of observations - whilst the process matches the variance profiles well, the persistence parameters and the variances for shocks are very noisily estimated.

The canonical process relies on the sequences of variances and autocovariances faced by each of the sub-cohorts that form a broad generation, and thus uses, for example, 122 observations. On the other hand, the NL

process relies directly on pairs of observations for earnings in t and $t + 1$, and uses 7500 such observations for the 1980s cohort.

Figure C.13 shows that, in the recalibrated model where earnings follow the canonical process, once households are out of their sizeable mortgages almost all of them invest in stocks, which does not happen in the baseline model

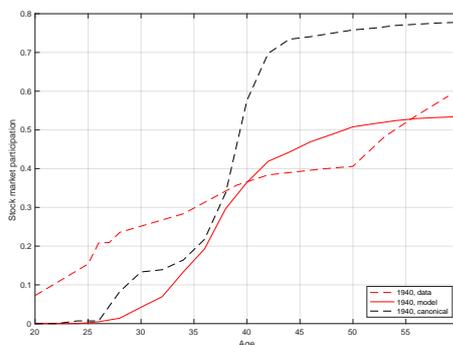


Figure C.13: Stock market participation for the 1940s cohort

3.5 Additional robustness checks

For all of this Appendix, unless otherwise stated, I assume that households start their working lives with zero wealth.

3.5.1 Initial zero wealth

Figure C.14 shows the fit of the model with respect to homeownership in the case in which I assume that all households are born with zero wealth and recalibrate parameters accordingly. Although the fit of homeownership at earlier ages is not as good, it is still true that the model can reconcile the changes in homeownership rates across generations without resorting to changes in preferences. Table C.2 shows that the decomposition of the decrease in homeownership between its different contributing factors is also similar to the main case. Finally, Figure C.15 shows that, without taking into account the initial condition, the model underpredicts stock market participation at earlier ages.

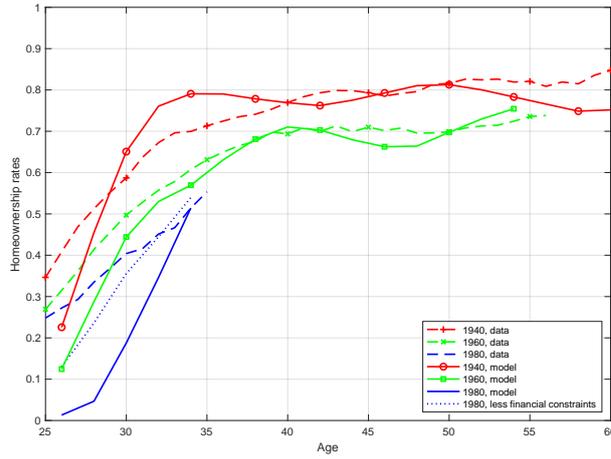


Figure C.14: Homeownership by cohorts, data vs. model, relaxation of financial constraints in the early 2000s. Model with zero initial wealth.

	1960s generation			1980s generation	
Age	30	40	50	30	35
Total	-9	-8	-9	-14	-22
Earnings	66	37	5	51	27
initial inequality	62	23	-17	33	8
risk	4	14	22	18	19
Aggregates	34	63	95	104	84
house price trend	73	129	63	90	74
histories	-39	-66	32	14	10
Financial conditions	0	0	0	-55	-11

Table C.2: Contribution of each factor in the change in homeownership with respect to the 1940s generation (% of the change), by age. Model with zero initial wealth.

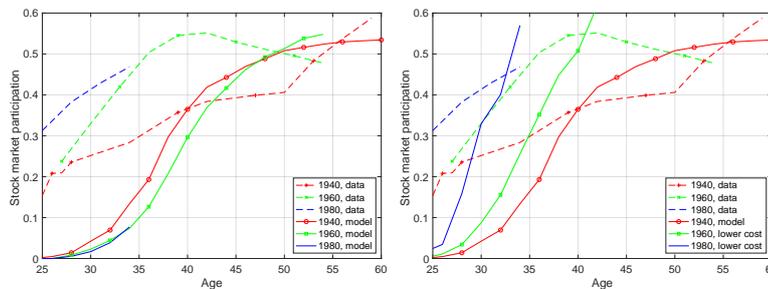


Figure C.15: Stock market participation by age and cohort, data vs model. Left: constant participation costs across generations; right: reduction of participation costs across generations. Model with zero initial wealth.

3.5.2 Empirical initial wealth and inter-vivos transfers

Although their net worth is usually very low, individuals between 20-25 years of age frequently have some wealth. Figure C.16, left panel, shows the empirical initial wealth distribution by cohorts. To the extent that this captures the average level of wealth of 20-25 year olds, it is also capturing any asset that comes from an early inheritance or family gift. However, this might underestimate inter-vivos transfers for two reasons. First, there might be family contributions (e.g. helping with the downpayment of a mortgage) which are not properly captured in the data. Second, these inheritances or family help might happen later than 25, and thus be difficult to separate in the data from wealth accumulation due to individual savings. In order to get an approximate idea of how far this channel could go in explaining the shortcomings of the model with zero initial wealth, I have used SCF data for the 1980s cohort to find out the percentage of people that got a large gift or inheritance before 35 (12.5%) and added its average value (one year of average income) to the top 12.5% of the initial wealth distribution.⁷⁵

Figure C.16, right panel, shows the results with the initial wealth distribution, and adding these additional inter-vivos transfers (IW+IV). Effects on homeownership profiles are minor when compared with the version of the model with zero initial wealth. However, the differences with the main version of the model, in which initial wealth positions are consistent with observed homeownership and stock market participation rates, suggest that initial wealth is poorly captured in these surveys, particularly for older generations.

⁷⁵This is an approximation, given that without a panel component it is not possible to find the joint distribution of (future) gift receipts and wealth.

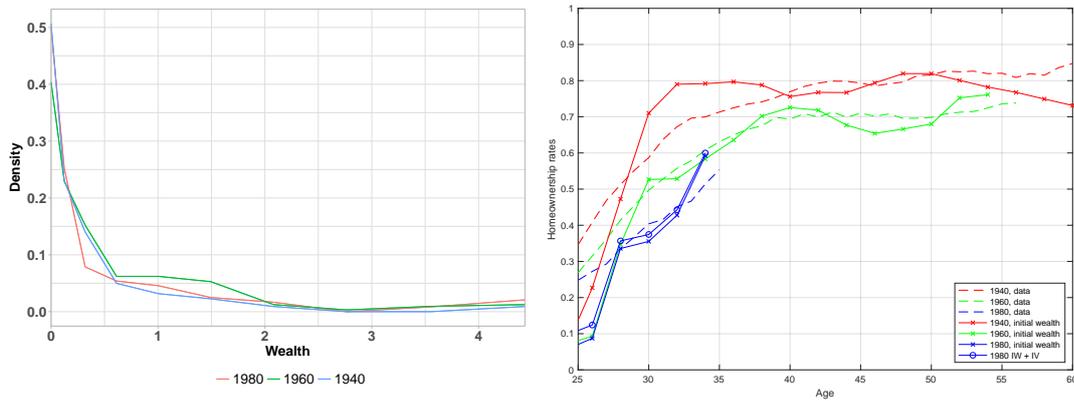


Figure C.16: Empirical initial wealth. Left: density of wealth of 20-25 year olds, SCF data, measured as multiples of average income. Zero includes zero or negative. Right: model-implied homeownership with empirical initial wealth

3.5.3 Equilibrium effects on stock returns

Figure C.17 shows to which extent the conclusions in Section 3.7.2 would vary as a result of possible general equilibrium effects on stock returns induced by the larger accumulation of financial wealth. In both cases, I assume that, once we enter the simulation period, yearly stock market returns fall unexpectedly and persistently by 2% or 4% for each possible realization of the stock state. A reduction of 2% in stock returns still implies that the 1980s cohort accumulates more financial wealth than that of the 1940s - a very significant reduction of 4% is required to make the profiles for both generations comparable. The accumulation of housing wealth is almost unaffected by these potential equilibrium effects.

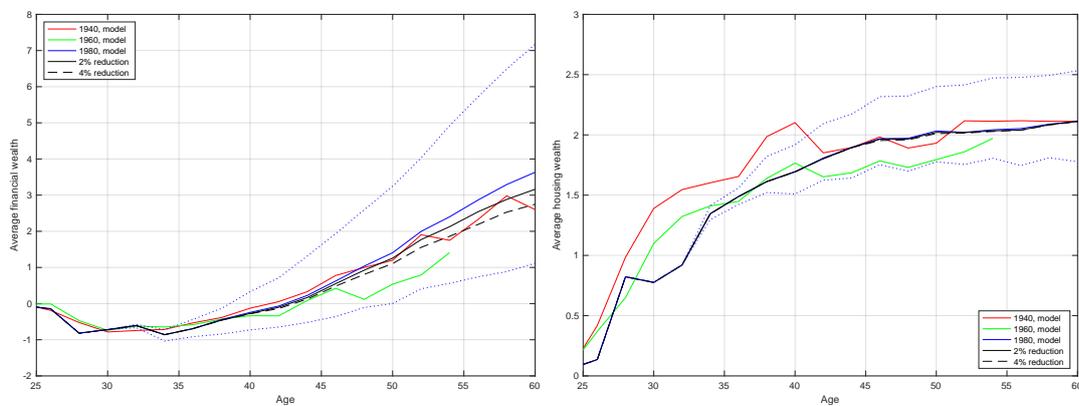


Figure C.17: Wealth accumulation: projecting the 1980s cohort into the future. With constant participation costs, comparing different reductions in average stock returns

3.5.4 Marital dynamics and family sizes

Over the past few decades, marriage rates have fallen, fertility has decreased, and the average age of women at both marriage and first child-bearing has steadily increased (Lundberg and Pollak, 2007). These developments offer a possible alternative explanation for the reduction and delay in homeownership.

In the framework proposed in this paper, all of those changes are implicitly considered in the earnings process, which is estimated in all households (married or not) and thus embeds marriage and divorce risk into earnings risk. In order to disentangle these two effects, the left panel Figure C.18 replicates the analysis by replacing the earnings process by one estimated only on continuously married couples. For this experiment, households start life with the same level of initial wealth as in the main version of the model. Given that this subset of households have on average higher and more stable earnings, their implied homeownership rates within the model are larger. However, differences across cohorts are present in a very similar way. This suggests that, while family dynamics can play a relevant role⁷⁶, the increase in inequality and earnings risk seems to be of first order to explain the changes in homeownership. To verify this hypothesis, the right hand side of Figure C.18 replicates the counterfactual in which the 1940s earnings process is attributed to younger generations within the subsample of married households. Effects are similar to those in the main results.

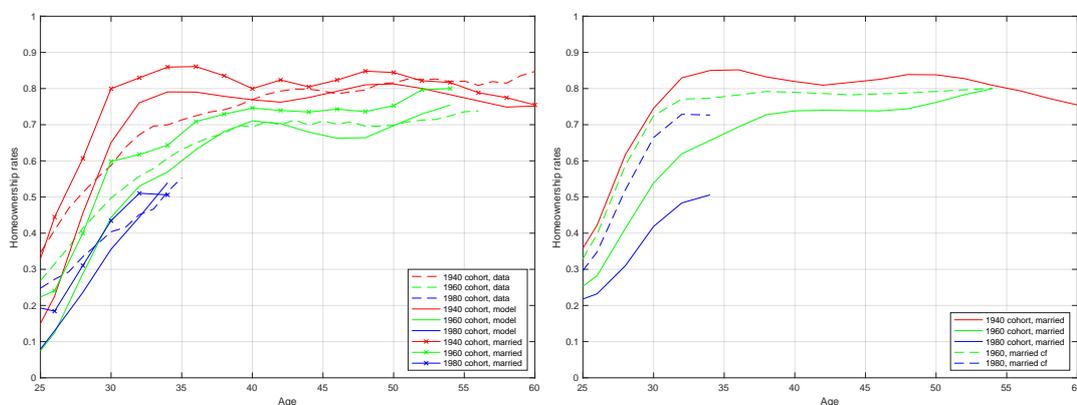


Figure C.18: Homeownership by cohorts, married people only.

⁷⁶Chang (2018) shows that marital and divorce risk is a relevant force to explain changes in homeownership of singles versus couples.

3.5.5 Per-period participation costs

There is a long standing discussion in the household finance literature about whether one-off entry costs, which I consider in the main version of this model, or per-period participation costs rationalize better the patterns of stock market participation that we observe in the data. While the former is used in studies like Cocco (2005) or Gomes and Michaelides (2005), Vissing-Jorgensen (2002), Gálvez (2017), or Bonaparte, Korniotis and Kumar (2018), amongst others, find that the latter seems to be the most promising avenue to explain the observed patterns of stock market participation. However, the estimated value for the per-period participation cost is frequently relatively high (for instance, Bonaparte et al. (2018) estimate it to be around 3.2% of average household income every year).

Figure C.19 shows the implied profiles for stock market participation, in a recalibrated version of the model that only allows for per-period participation costs that decrease over time. As in previous studies, the estimated cost which is necessary to reconcile low levels of stock market participation is quite high (around 3.5% of average household income for the 1940s cohort), but is lower for younger cohorts.

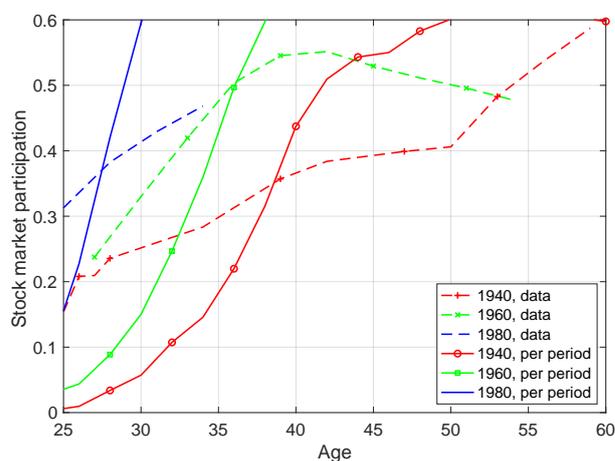


Figure C.19: Stock market participation over the life-cycle, by cohorts, case with calibrated per-period participation costs

3.5.6 Local correlation of income shocks and house prices

An additional element that increases the riskiness of housing is the correlation at the local level between house prices and income changes. For

instance, in areas that benefit in particular from an expansion it is likely that both incomes and house prices go up. When households incorporate this information into their decision making, it influences both homeownership and portfolio choices.

To approximate this effect, I allow for income shocks to be correlated with housing price shocks at the idiosyncratic level. Given that in the baseline version of the model both are exogenous, this correlation can be directly imposed from data estimates. I rely on [Davidoff \(2006\)](#) for the empirical quantification of this correlation and fix it at 0.29.

Figure C.20 shows the corresponding homeownership profiles. I perform two experiments. The left hand side panel represents the case in which house prices are expected to be correlated with income, and in which the realizations of house prices are also correlated with the realizations of the income shocks. This induces, unlike in the main case in this paper, housing price heterogeneity within agents living in the same year. Like in the main case, I assume that the average realization of the house price shock is like that observed in the data for that specific year. Homeownership would be lower in this case for all cohorts, particularly for those born in the 1980s, even with looser financial constraints.

The right hand side panel represents the alternative case in which households expect house prices to be correlated with their individual income shock, but where this is not true in realization and house prices are still their corresponding national average. Differences with the previous experiment are minor, which suggests that the key driver for this reduction in homeownership rates are household portfolio decisions rather than changes in the stationary distribution induced by the introduction of house price heterogeneity.

This experiment abstracts from important elements such as endogenous determination of local house prices or endogenous mobility, which are beyond the scope of this project and are left for future research.

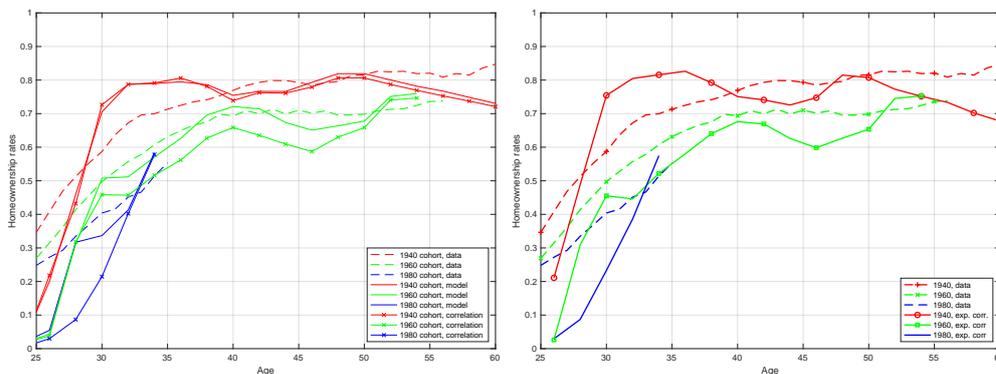


Figure C.20: Homeownership by cohorts, local correlation of house price shocks with income. Solid line with crosses: realizations + expectation, dotted line: expectations only

3.5.7 House sizes

In the main version of the model, households can choose to own two types of housing, which I denote h_1 and h_2 . This reflects that houses are lumpy and have limited divisibility, and that frequently households cannot access their optimal house size and quality because it is scarcely available or disadvantageously priced on the market. In this section, I check the robustness of my results with respect to different specifications of house sizes.

The size of the small house h_1 is key for my results, as it conditions who is the marginal person who is indifferent between buying a house and renting, and thus homeownership rates. In current U.S. dollars, for the latest periods in the model, the price of a small house is around \$100,000. Here I argue that there is limited supply of houses below this price, few households actually live in them, and including a smaller small house in the model generates counterfactual implications.

With respect to supply limitations, [Falcettoni and Schmitz \(2018\)](#) argue that regulations and monopoly power have reduced the production of prefab or factory-built houses below its efficient level, thus reducing the number of cheap houses available on the market. In my PSID data, I find that, for the 1980s cohort, only 9.4% of households bought a house with a price lower than h_1 , with most of them concentrated nearby.

The left hand side panel of Figure C.21 shows the model implications for the case in which the smallest house is 33% smaller and cheaper. Nat-

urally, results are sensitive to this assumption. However, as Figure C.22 reveals, this experiment underestimates generational differences by overestimating how many people buy small houses as a reaction to changes in earnings dynamics and housing prices.

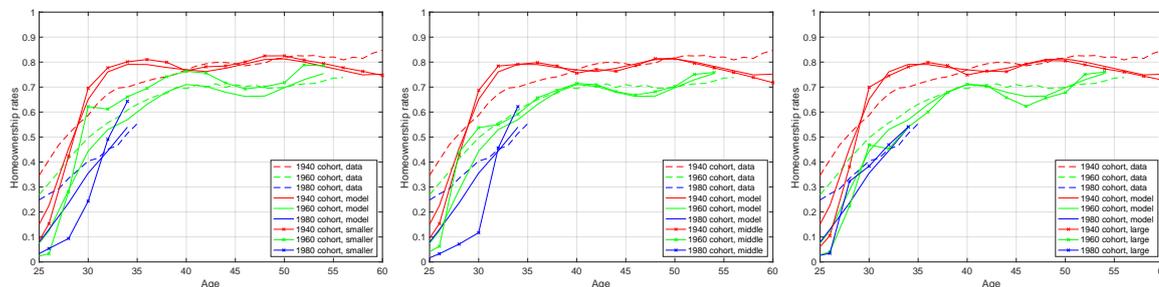


Figure C.21: Alternative house sizes, main experiments

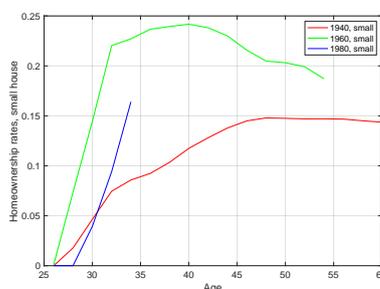


Figure C.22: Percentage of households owning the small house by age. Left, robustness check in which it's smaller; right: main results

The central and right panels of Figure C.21 show two robustness experiments in which I set the number of housing qualities to $H = 3$. In the middle panel, I do so by allowing for a middle-sized house exactly between the prices of the small house and the big house. In the right hand side panel, I allow for a large house, twice the size of the big one. In all cases I recalibrate housing preference parameters accordingly so that homeownership for the 1940s cohort at age 40 is within the ballpark of the data.

Results are almost identical with either of the $H = 3$ assumptions, thus suggesting that the choice of $H = 2$ is not key in driving the results I obtain. Setting $H = 3$ and allowing for a large house does improve the fit of the model in terms of the portfolio composition of the richest, which slightly underestimates their housing share (see Figure 32).