



**Trinity
College
Dublin**

The University of Dublin

Essays in Macroeconomics and Wealth Inequality

Author: Simone Cima

Supervisor: Prof. Joseph Kopecky

March 2025

*A thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy
in the Department of Economics*

DECLARATION

I hereby declare that this thesis has not been submitted as an exercise for a degree at this or any other university. All research contained herein is entirely my and my co-authors' work and all errors are our own. Chapter 2 is co-authored with Joseph Kopecky, and Chapter 3 is co-authored with Marco Moreno. All the views expressed in this thesis are mine and my coauthors', and do not necessarily represent the views of the Central Bank of Ireland or the Eurosystem.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

Signed: _____ Date: _____

Summary

This dissertation consists of three essays studying the dynamics of wealth inequality as it relates to different macroeconomic developments. In the three essays, I employ a range of different methods used in macroeconomics, and both empirical and theoretical frameworks. Each of the essays contributes a novel perspective, whether through the specific focus, or the use of data, to its respective literature.

Chapter 1 studies how the shape of an economy's population distribution, or population pyramid, can affect rates of return and incomes. In turn, I observe how these developments can lead to unequal wealth accumulation journeys of cohorts of the population born at different points in time. Because older populations accumulate more wealth and have lower employment rates, young people in an ageing society face lower returns, and as a result accumulate wealth more slowly; at the same rates of savings, this implies lower standards of living at retirement for these cohorts compared to the cohorts who were of working age during population booms, leading to inter-generational inequality. I employ instrumental-variable fixed-effect panel regressions, using a unique parameterisation of the shape of the population distribution, and a number of back of the envelope calculations.

Chapter 2 analyses the effect of a number of housing market policies on inequality in housing wealth, wealth, income and consumption, as well as homeownership. This is done by constructing a heterogeneous agent overlapping-generations model which features endogenous choice between renting and homeownership, and both a housing and a rental market. Within this framework, we study the effect of macroprudential policies, rental income taxation, the role of institutional investors in the housing market, supply side housing policies, and interest rates, as well as combinations thereof. We highlight a number of insights from this exercise, with sometimes unintuitive results, that can be of great value

for policymakers. We also provide a solid framework for future work exploring the effects of the same policies with more complexities, as well as other policies.

Chapter 3 estimates the differential effect of ECB monetary policy shocks on the wealth of euro area households across the wealth distribution. We employ panel local projections on newly published, quarterly distributional wealth data for the euro area, and identified monetary policy shocks. We find that the wealth levels and compositions of households at different positions in the wealth distribution respond differently. For instance, the net wealth, asset and debt levels of the bottom half decline significantly in response to a contractionary monetary policy shock, while the wealth of the top 10% of the distribution shows an almost opposite behaviour. We also highlight significant portfolio rebalancing across the distribution, and an asymmetry in the responses to contractionary versus expansionary shocks.

Acknowledgements

This thesis would not have been possible without the invaluable help and support of a large number of people. First of all, I would like to express my greatest gratitude to my supervisor, Professor Joseph Kopecky, who is also coauthor of the second chapter of this thesis. His invaluable advice, knowledge, friendly disposition and constant encouragement have been indispensable. Thanks most of all for helping me grow as an economist and researcher.

I would also like to thank all faculty and staff of the Economics school in Trinity. Thanks also to my fellow PhD students at TCD: while doing this part-time I have had fewer occasions than I would have liked to engage with you, you have made Trinity such a welcoming, refreshing and constructive environment.

I am grateful to Michael Wycherley, Davide Romelli, Paul Scanlon, Reamonn Lydon, Gerard O'Reilly, Fang Yao, Fergal McCann, Gabriel Arce-Alfaro, Stephen Byrne, Giuseppe Corbisiero, Tara McIndoe-Calder, Laura Boyd, Simone Arrigoni, Alexis Lager, Vahagn Galstyan, Boryana Madzharova, and others for helpful and insightful comments and suggestions I received from them on the various chapters of this thesis. My thanks also go to Vitaliia Yaremko and Henrique Basso, for a constructive discussion of my thesis and for their valuable feedback.

Thanks to the Central Bank of Ireland for providing the funding and support I have needed throughout this journey, and to all my colleagues, past and present, for their advice and suggestions, but most of all for contributing to making work and research such an enjoyable and fulfilling part of my life. You are too many, but I would like to single out Neill Killeen, for making this PhD possible in the first place and encouraging me to do it, and Stephen Byrne and Reamonn Lydon, for their extraordinary support and for making me grow professionally. To my colleague (and coauthor of the third chapter of this thesis), Marco Moreno, thanks for

being the ideal version of both.

My heartfelt thanks to all those friends, in Dublin, Italy and elsewhere (you know yourselves!), who have supported me in this journey and cheered me on. Elena Sponga, simply thanks for existing and being the amazing friend you are.

Thanks to my honorary research assistant, my dog Parmigiano, who kept me company through much of my struggles with coding and the literature. He clearly cares about inequality, too, especially when I do not share food with him.

To 森, for being by my side this whole time, through the highs and lows. I am looking forward to whatever comes next, together.

Thanks to my family, for always having my back. I would be nowhere without you. This is for my mum.

Contents

Declaration	i
Summary	iii
Acknowledgements	v
Contents	vi
List of Figures	ix
List of Tables	xiii
General Introduction	xv
1 An inverted-pyramid scheme? Demographics, wealth accumulation and inter-generational inequality	1
1.1 Introduction	1
1.2 Literature Review	4
1.3 Data	7
1.3.1 Sources	7
1.3.2 Total Return on Capital	8
1.3.3 Population and Demographic Distribution	12
1.4 Demographics and Rates of Return	14
1.4.1 Working-age Population and Productive Capacity . .	14
1.4.2 Age Distribution and Capital Accumulation	19
1.4.3 Returns to Wealth and Inequality Between Generations	25
1.5 Demographics and Wages	33
1.6 Wealth Accumulation: Winners and Losers	38
1.7 Conclusions	46
2 Housing policies, homeownership and inequality	50
2.1 Introduction	50

2.2	Model	57
2.2.1	Households	57
2.2.2	Government	60
2.2.3	Production	61
2.2.4	Institutional Investors	62
2.3	Solution and calibration	64
2.3.1	Calibration	65
2.4	Results	70
2.4.1	The baseline model	73
2.4.2	Tightening of macroprudential measures	74
2.4.3	Institutional investors	78
2.4.4	Tax on rental income	84
2.4.5	Measures on the construction sector	86
2.4.6	A not implausible negative scenario	89
2.5	Conclusions	90
3	Monetary policy shocks across the wealth distribution: insights from new European data	92
3.1	Introduction	92
3.2	Data	97
3.2.1	Composition, distribution and dynamics of wealth in the euro area	99
3.3	The empirical approach	103
3.4	Monetary policy and wealth dynamics across the wealth distribution	107
3.4.1	Levels	107
3.4.2	Portfolio composition	110
3.4.3	Wealth inequality	111
3.5	Conclusions	114
	Bibliography	116
A	Appendices	125
A	Measures of productivity	125
B	Comparison of the various simulations	128
C	Changes in interest rates	131
D	Results over tenancy and working status breakdowns	133
D.1	Homeownership breakdown	133
D.2	Employment breakdown	136
E	Additional charts	140

List of Figures

1.1	<i>Evolution of the Japanese population pyramid.</i>	2
1.2	<i>Average Real Capital Price Index, and Components</i>	11
1.3	<i>Example of a demographic distribution and its polynomial representation based on the D_1, D_2 and D_3 variables. USA, 2021 (as from UN World population prospects 2019).</i>	14
1.4	<i>Effect of each age group on $\ln(K/N)$ and $\ln(L/N)$</i>	22
1.5	<i>Combined effect of each age group on rates of return via $\ln(K/N)$ and $\ln(L/N)$</i>	24
1.6	<i>Impact of the demographic distribution on return on capital (r^k), selected advanced economies.</i>	27
1.7	<i>Returns to capital and demographic effect: CHN, ITA, JPN, USA</i>	28
1.8	<i>Wealth accumulated over 20 years: CHN, ITA, JPN, USA</i>	31
1.9	<i>Combined effect of each age group on compensation of labour via $\ln(K/N)$ and $\ln(L/N)$</i>	36
1.10	<i>Impact of the demographic distribution on compensation of labour (w^L), selected advanced economies.</i>	37
1.11	<i>Evolution of (log) real wages per person employed (actual versus with demographic effects removed), selected advanced economies.</i>	38
1.12	<i>Wealth accumulated over 20 years from initial investment of 10% of wage, selected advanced economies.</i>	41
1.13	<i>Wealth accumulated over 20 years from initial investment of 10% of wage, as a fraction of peak value, average across selected countries.</i>	42
1.14	<i>Wealth accumulated over 40 years' working life, as a fraction of peak value, average across selected countries.</i>	43
1.15	<i>Annual cash flow from reverse mortgage of wealth accumulated over 40 years' working life, as a fraction of peak value, average across selected countries.</i>	45
1.16	<i>Annual cash flow from reverse mortgage of wealth accumulated over 40 years' working life, as a fraction of peak value.</i>	46

2.1	<i>Illustration of multiple equilibria in the housing market when supply is irrelevant.</i>	71
2.2	<i>Surface area representation of demand for housing as a percentage of population under the baseline model, by p_h and r_h.</i>	72
2.3	<i>Gini coefficients for wealth, disposable income and consumption, for each (p_h, r_h) combination leading to equilibrium ($H_d = N_t$) in the baseline model.</i>	73
2.4	<i>The construction sector determines the optimal price, while intersection with housing demand determines the equilibrium rental rate.</i>	74
2.5	<i>Effect of a tightening of macroprudential policy forcing some households into the rental market. Seen in two different ways.</i>	75
2.6	<i>The presence of institutional investors pins the rental return. This might reduce household sector home-ownership and wealth levels at the top, but also rents.</i>	80
2.7	<i>Impact of the presence of institutional investors in the housing market on the distribution of housing wealth, consumption and total wealth, with LTI of 6 and r^{ii} of 1/3%.</i>	81
2.8	<i>A tightening of macroprudential policies when institutional investors are present lowers the number of housing units owned by domestic households (whether they are an investment or owner-occupied).</i>	83
2.9	<i>Lowering τ_{rh} benefits both renters and landlords, but not in equal measure</i>	85
2.10	<i>Expanding construction permits leads to a lower p_h allowing for $H_s = N_t$. This in turn leads to a lower equilibrium r_h</i>	87
3.1	<i>Composition of household portfolios in the euro area as a percentage of total assets, by net wealth group.</i>	100
3.2	<i>Composition of household portfolios in the euro area, by net wealth group (millions of 2015 €).</i>	102
3.3	<i>Household group's share of total holdings in each asset type, euro area.</i>	104
3.4	<i>IRFs: cumulative percent changes in households' non-risky financial assets following a one-standard-deviation contractionary monetary policy shock, by net wealth group</i>	108
3.5	<i>IRFs: cumulative percent changes in households' housing assets following a one-standard-deviation contractionary monetary policy shock, by net wealth group</i>	109

3.6	<i>IRFs: cumulative percent changes in households' net wealth following a one-standard-deviation contractionary monetary policy shock, by net wealth group</i>	110
3.7	<i>IRFs: change in the share of non-risky financial assets in households' portfolios following a one-standard-deviation contractionary monetary policy shock, by net wealth group</i>	111
3.8	<i>IRFs: household group's share of country's net wealth following a one-standard-deviation contractionary monetary policy shock . .</i>	112
3.9	<i>IRFs: household group's share of country's non-risky financial assets following a one-standard-deviation contractionary monetary policy shock</i>	113
3.10	<i>IRFs: household group's share of country's net wealth following a one-standard-deviation contractionary monetary policy shock (blue lines) vs a one-standard deviation expansionary monetary policy shock (red lines)</i>	114
A.1	<i>Combined effect of each age group on rates of return via $\ln(K/N)$ and $\ln(L/N)$, by measure of productivity used in instrumental variable regression (confront Tables 1.2 and 1.3).</i>	126
C.1	<i>Surface area representation of demand for housing as a percentage of population, when r_b is increased to 1%, all else equal, by p_h and r_h.</i>	132
D.1	<i>IRFs: percent changes in households' net wealth following a one-standard-deviation contractionary monetary policy shock, by tenancy status</i>	134
D.2	<i>IRFs: household group's share of country's net wealth following a one-standard-deviation contractionary monetary policy shock, by tenancy status</i>	135
D.3	<i>IRFs: household portfolio share of non-risky financial assets following a one-standard-deviation contractionary monetary policy shock, by tenancy status</i>	135
D.4	<i>IRFs: percent changes in households' net wealth following a one-standard-deviation contractionary monetary policy shock, by employment status</i>	137
D.5	<i>IRFs: households' portfolio share of housing assets following a one-standard-deviation contractionary monetary policy shock, by employment status</i>	137

D.6	<i>IRFs: households' country share of net wealth following a one-standard-deviation contractionary monetary policy shock, by employment status</i>	138
E.1	<i>Composition of household portfolios (detailed breakdown) in the euro area, by net wealth decile.</i>	140
E.2	<i>Composition of household portfolios (detailed breakdown) in the euro area, by net wealth decile.</i>	141
E.3	<i>IRFs: per cent changes in households' debt liabilities following a one-standard-deviation contractionary monetary policy shock, by net wealth decile</i>	142
E.4	<i>IRFs: per cent changes in households' risky financial assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile</i>	142
E.5	<i>IRFs: per cent changes in households' business assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile</i>	143
E.6	<i>IRFs: households' debt over total assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile</i>	143
E.7	<i>IRFs: households' portfolio share of housing assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile</i>	144
E.8	<i>IRFs: households' portfolio share of risky assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile</i>	144
E.9	<i>IRFs: households' portfolio share of business assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile</i>	145
E.10	<i>IRFs: household group's share of country's debt liabilities following a one-standard-deviation contractionary monetary policy shock</i>	145
E.11	<i>IRFs: household group's share of country's non-risky financial assets following a one-standard-deviation contractionary monetary policy shock</i>	146
E.12	<i>IRFs: household group's share of country's risky financial assets following a one-standard-deviation contractionary monetary policy shock</i>	146
E.13	<i>IRFs: household group's share of country's business assets following a one-standard-deviation contractionary monetary policy shock</i>	147

List of Tables

1.1	Fixed-effects panel estimation of return on capital based on equations 1.4 and 1.5.	17
1.2	Instrumental-variables estimation of r^k . First-stage regressions.	21
1.3	Instrumental-variables estimation of r^k . Second-stage regression.	23
1.4	Instrumental-variables estimation of real (log) w^L . Second-stage regression.	35
2.1	Calibrated parameters of the baseline model.	66
3.1	Breakdown of households net wealth items in the DWA . .	98
A.1	Specification (2) of Table 1.1, for different measures of productivity.	127
B.1	Model outcomes, various model specifications	129

General Introduction

"Quem tenha lido as páginas deste livro, que estão antes desta, terá sem dúvida formado a ideia de que sou um sonhador. Ter-se-ia enganado se a formou. Para ser sonhador falta-me o dinheiro."

- Fernando Pessoa, *Livro do Desassossego*¹

In an unequal society, many people cannot afford to be dreamers, but they should be able to. Inequality can be very damaging not only for the individuals most affected by it, but also for the economy as a whole. There are many causes for it, and as many effects of it. With this thesis, I hope to shed light on a few of these.

Distributional issues have in the past often been typically overlooked in economics, and especially macroeconomics, the main focus of studies usually resting on aggregates and their dynamics. More recently, however, the discipline has developed a deep understanding of how not only do aggregate movements matter for distributional outcomes, but also how, conversely, the distributions of income and wealth themselves matter for aggregate outcomes, and how we cannot fully comprehend a phenomenon without understanding how it affects (and is affected by) the behaviour of households in different circumstances. But inequality is not only an important issue as it relates to other aspects of the economy: to understand it and reduce it is a fundamental goal of economic policy in itself. No society can be truly called successful where a portion of the population lives precariously or in poverty.

In this thesis, I study different aspects and sources of wealth inequality, utilising different techniques, and showing how very different aspects of our

¹This work is under public domain. Loosely self-translated as: "Whoever has read the above pages of this book, would have doubtlessly got the idea that I am a dreamer. They would be mistaken. I am lacking the money to be a dreamer."

economies, from the shape of the population distribution, to policies aimed at the housing market, to monetary policy, can have large distributional implications. Indeed, whether temporarily or permanently, households are never all affected equally by any dynamics in the macroeconomy; this, one way or another, leads to developments in inequality. Knowing the causes and dynamics of these developments is a fundamental step in trying to address inequality and/or mitigate its effects.

This thesis consists of three essays: in Chapter 1, I look at how a country's demographic distribution affects returns on wealth and wages, and I examine the implications of this for different generations as populations age; in Chapter 2, co-authored with my PhD supervisor Joseph Kopecky, we estimate the implications of a number of housing market policies on homeownership and inequality of wealth, income and consumption, within an overlapping generations framework; finally, in Chapter 3, co-authored with Marco Moreno, we study the heterogeneous responses of euro area households' wealth, portfolio allocations, and investment decisions to ECB monetary policy shocks, using a novel dataset.

In the first Chapter, I study how inequality in the standards of living of different generations can arise in ageing societies going through a demographic transition. I observe how demographic dynamics and the shift of population pyramids towards an inverted-pyramid shape in advanced economies are leading to relative scarcity of labour and excess savings. I explore the effects of these dynamics on the relative wealth accumulation journeys of population cohorts born in different generations. Within a fixed-effect cross-country panel framework, I find that economies with ageing populations are associated with large levels of accumulated wealth per capita, which in turns makes it harder for younger cohorts to accumulate wealth for themselves. Savings by an increasing share of households aged between 45 and 65, a rise in retired over-65s, and a decrease in working-age and low-wealth agents in their twenties and thirties can explain for a large part of the decline in rates of return across countries in the last few decades, and similarly a large part of the increase in wages.

Within this context and looking to the future, plausible forecasts suggest that, due to demographic developments, wealth accumulation out of income and capital returns by population cohorts currently living and in work in advanced economies is set to become increasingly difficult in the next decades. Unless they receive conspicuous inheritances, future retirees in these cohorts are thus set to face progressively lower standards of living

at retirement and/or increasingly high saving ratios in working age. This occurs as projected higher wages are not sufficient to compensate for the lower returns on capital that they are going to be facing over long periods of time. These cohorts are therefore likely going to be worse-off compared to current and past retirees, once they retire. That is even more the case if we also consider rising life expectancies and the need to finance longer retirements.

The second Chapter of this thesis is interested in quantifying how a number of different policies in (or features of) the housing market can affect measures of inequality, household welfare through the population, and homeownership. To this end, we develop a heterogeneous agents overlapping-generations life-cycle model with housing and rental markets. We focus on macroprudential policies, institutional investors, taxation of rental income, and supply-side policies focused on the construction sector.

By restricting credit to lower income households, macroprudential regulation has a negative impact on homeownership rates while it also causes increases in rental yields, and as a result it leads to higher inequality. Supply-side housing policy aimed at expanding housing supply can reduce both house prices and rental yields, benefiting every household except landlords and greatly reducing post-housing income and consumption inequality; these policies actually do not necessarily give rise to higher homeownership rates.

We further explore the effect of the introduction of institutional investors into the rental market, and find that they can decrease inequality by lowering equilibrium rental rates, at the cost of reduced homeownership. Taxation on rental income, targeting the incomes of the wealthiest, might actually harm renters the most by pushing up rental payments. As different policies can lead to different trade-offs between different types of inequality, as well as homeownership rates, we note how an optimal policy mix can be achieved by combining these policies in an effective way. Our model and framework are suited to a wide range of further avenues for exploration in this area.

In the third Chapter of this thesis, we use new data on the distribution of wealth in the euro area, the Distributional Wealth Accounts, and employ panel local projections to estimate the differential impact of ECB monetary policy shocks on euro area households across the wealth distribution. We

look at how policy affects their assets, the composition of their balance sheets, and overall wealth inequality. We reach results that are broadly in agreement with the literature on the effects of monetary policy on inequality, while providing a new empirical angle, with distributional data at a quarterly frequency, to the discussion.

We find that, in response to a contractionary monetary policy shock, financial assets of most households are hit strongly, with housing less affected. It is poorer households in particular, however, who display a substantial decline in their assets. Conversely, the balance sheet of the very wealthiest shows the opposite evolution, as households in the top 10% of the wealth distribution see their net wealth increase in response to the same shock. This ultimately leads to an increase in overall wealth inequality, as top wealth holders end up owning a larger share of wealth in most asset classes. We note that these different dynamics cannot be due exclusively due to movements in asset prices, but also different behaviour of households in different situations. In terms of households' balance sheet composition, we find that contractionary monetary policy shocks shift portfolios away from financial assets; this appears to happen across the whole wealth distribution, apart from the very top. We also find that expansionary shocks, while generally with opposite sign to contractionary shocks, are not exactly symmetrical to them.

1 An inverted-pyramid scheme? Demographics, wealth accumulation and inter-generational inequality


1.1. INTRODUCTION

The world is in the midst of one of the most dramatic demographic shifts in history. After two centuries of exponential expansion, world population growth reached its peak in the 70s and has experienced a considerable slowdown since. The ageing of societies is already at an advanced stage in many developed economies, some of which face population declines. This is quickly becoming a global phenomenon. Life expectancies have never been as high; meanwhile, birth rates are well below replacement in an increasing number of economies.

Population pyramids in a number of countries have now started taking the shape of *inverted* pyramids, with younger generations progressively smaller and a high median age - see Figure 1.1 for the striking example of Japan. Population structures of this type are unprecedented. An ageing and eventually declining population has profound consequences for the prospects of an economy, in terms of productivity, equality, allocation of resources, financial markets and economic policy.

As longer lifespans push larger shares of workers towards retirement age, while the relative size of the working population shrinks, an obvious question arises: how to finance the livelihoods of current retirees while allowing the working-age population to accumulate enough wealth to fund their own retirement? With public finances increasingly strained by pension payments, a large part of retirees' income might need to originate from the

Figure 1.1: *Evolution of the Japanese population pyramid.*



Percentage of the population in each 5-year age group.

private wealth accumulated during their working life. The accumulation of sufficient amounts of wealth requires not only a steady stream of savings, but also high enough real rates of return on such savings. However, population aging itself might have an adverse effect on rates of return on wealth; empirical results obtained in this paper give further evidence of this.

In the context of demographic ageing and decline, younger generations in an increasing number of countries have been worrying that their future economic prospects look rather bleak when compared to their parents' generations. When their parents could, for instance, afford a house and start a family early in life, this is becoming increasingly difficult for many young people in advanced economies, who are facing a situation of low savings with low returns, and high asset prices. Are these worries founded, and do demographic dynamics have anything to do with trends in wealth inequality?

The main purpose of this paper is to explore and quantify how the wealth accumulation journey of different generations (or cohorts) has been or will be affected by demographic trends, and to derive implications in

terms of inequality (of wealth, but more generally, of standards of living) between these cohorts. While the analysis within the paper explores these dynamics across the world, particular emphasis is given to older, advanced economies at the frontier of the demographic transition.

In these economies, how do young workers today fare compared to their parents' generation? Generally speaking, will they be worse-off? I find that, in advanced, ageing economies, that appears indeed to be the case. Consistent with a growing literature, this paper presents evidence that ageing pushes down rates of return, but also positively affects wages. As these dynamics unfold, accumulating wealth becomes relatively more difficult for at least some section of the population. Because returns will be lower, any initial amount of saving will not accumulate as quickly. At the same time, wages will increase, but not enough for households to be able to reach the same levels of wealth accumulation as the generations before them with the same saving rates, unless unlikely high productivity growth is assumed in future decades. Younger generations will thus have to save more to reach the same levels of wealth as their parents, and even more to reach a level of wealth that will allow them to have an equally comfortable retirement (seeing as life expectancies are set to increase further).

I reach these findings by employing a simple fixed-effect panel estimation approach. I find that ageing demographics tend to depress rates of return and increase wages, although with different intensities, through the effect of the population structure on capital accumulation in the economy and employment. The estimation method used also allows me to show how each different age group has a markedly different impact on returns and wages through these two channels. By combining the single age-group effects together, I can then recover how the whole population pyramid affects rates of return and wages, and wealth accumulation as a result, throughout countries and years. With a number of quantitative exercises and back of the envelope calculations, I then finally quantify the implications of these dynamics for the wealth accumulation journeys of different cohorts, and how they fare compared to one another, reaching the conclusion that current young generations and those in the future are set to lose, especially compared to the generation currently at retirement age, which in the average advanced economy benefited from the most favourable combination of demographic and economic conditions.

The rest of the paper is structured as follows: Section 1.2 reviews some of the literature on demographics, wealth accumulation, the macroecon-

omy and inequality, and locates this paper in the context of this literature; Section 3.2 introduces the data used in the paper and addresses some issues related to these data; Section 1.4 looks at the impact of demographic dynamics on rates of return through different channels and explores the implications for wealth accumulation and inter-generational inequality; Section 1.5 repeats this analysis for wages; Section 1.6 combines these findings to obtain the overall impact of demographics on wealth accumulation across generations; finally, Section 3.5 concludes the paper.

1.2. LITERATURE REVIEW

This paper draws from a vast and established literature on demographics and the economy, of which I can identify a few strands that are most relevant in particular to the effects of demographic dynamics on rates of return, and inequality.

One of the main channels through which demographic developments affect the economy is through savings. The literature on life-cycle consumption and savings generally finds that the consumption and saving behaviour of economic agents changes as they age, as their earnings and levels of wealth change, and they prepare for retirement (Modigliani, 1988). Households start their working life with low earnings and quickly increase their incomes, and earn most close but before retirement age, after which their income falls. Consumption is also hump-shaped and partly follows the evolution of income, however saving rises by more for several decades before declining sharply and becoming negative well into retirement. The literature on life-cycle behaviour is vast and established - see for instance Gourinchas and Parker (2002) for a life-cycle representation of income and consumption. Results in Feiveson and Sabelhaus (2019) show that annual saving is typically highest when an agent is between 35 and 50, while wealth increases at the fastest pace between 40 and 60 (due to effects of returns on accumulated wealth). Aggregating this type of saving and consumption behaviour in the macroeconomy results in large effects of age on aggregate wealth, consumption, and saving rates. In particular, a large share of the population in the ages with the highest accumulated lifetime wealth will lead to higher average levels of wealth per person in the economy, with implications for interest rates.

There is ample evidence, both theoretical and empirical, that population ageing indeed leads to a reduction in interest rates - whether by interest

rate we intend the natural interest rate or rates of return or other measures - although there is disagreement on the magnitude and on how other effects are also at play. A large number of papers reach these conclusions by employing an overlapping-generations (OLG) model; these for instance include [Geppert et al. \(2016\)](#), [Börsch-Supan et al. \(2003\)](#), [Lisack et al. \(2017\)](#), and [Papetti \(2021\)](#); [Krueger and Ludwig \(2007\)](#) also find that, in an open economy, capital flows cause economies to “export” or “import” part of the decline in interest rates due to demographics (respectively, if ageing happens more or less rapidly compared to ROW). [Carvalho et al. \(2016\)](#) again find, within a life-cycle model, that ageing demographics lead to a reduction in interest rates, via increasing capital per worker and higher life expectancies, although lower savings by retirees have an opposite but insufficient effect (similar results are found in [Papetti \(2021\)](#) and [Börsch-Supan et al. \(2003\)](#)). [Ikeda and Saito \(2014\)](#) attribute a greater part of the reduction in interest rates in Japan to a decline in TFP, but demographics also have a large negative effect. That ageing demographics put downward pressure on interest rates is also documented empirically - see for instance [Ferrero et al. \(2019\)](#).

A reduction in interest rates, generally, tends to be a collateral finding of a great part of theoretical papers with demographic dynamics. These multiple results are all consistent with the literature on the life-cycle patterns of savings and consumption, as in general a higher proportion of the population in ages with high accumulated savings leads to an excess of these relative to investment, thus lowering rates. Some papers however challenge the view that demographics are the main factor responsible for decreasing rates. [Goodhart and Pradhan \(2017\)](#) mostly attribute falling interest rates in the past decades to a shift of manufacturing production to emerging markets (especially China), and rather see ageing demographics globally as potentially reversing the trend in the near future. [Mian et al. \(2021\)](#) mostly associate the fall in the natural rate of interest in the US to increases in income inequality leading to higher saving by the wealthier.

In this paper, I find again that ageing demographic structures lead to lower rates of return on capital through higher capital accumulation and lower employment. Interestingly however, thanks to the estimation I employ, I am also able to go further and break down the effects on rates of return by age group and to quantify how each specific shape of the population pyramid might impact returns. The results suggest that demographic dynamics may account for most of the variation in rates of return across

the years in the countries studied.

Demographic ageing is not only associated with lower interest rates, but also with lower potential growth - papers that make this argument include [Cooley and Henriksen \(2018\)](#) and [Aksoy et al. \(2019\)](#). Within a context of population ageing and low or negative population growth, combined with anemic productivity growth and low interest rates (which due to demographic dynamics are all seen occurring together and persisting for the rest of the century) part of the literature talks about secular stagnation. Demographics are only one of the potential explanations for the low interest rate environment in advanced economies and secular stagnation - both [Sajedi and Thwaites \(2016\)](#), and [Eichengreen \(2015\)](#), for instance, make the case for the fall in the relative price of investment goods being one contributing factor. In his famous speech, [Summers \(2014\)](#) lists six structural reasons for the decline in natural interest rates (which he links to secular stagnation) experienced in the US, and out of these, one is the slowdown in population growth. While in this paper I do not make considerations on economic growth and productivity as it relates to demographics, this strand of the literature informs the analysis and interpretation of results.

With life-cycle savings, rates of return and employment as main channels, the literature tends to find that ageing demographics lead to increasing wealth inequality, particularly of the inter-generational type, although there is no general consensus. [Antunes and Ercolani \(2020\)](#) find that demographic dynamics can explain a large part of the increase in inter-generational wealth inequality in the US. [Geppert \(2015\)](#) finds that demographic effects on wages, human capital accumulation and the skill premium will lead to different welfare effects across household types. [Vandenbroucke \(2016\)](#) sees increases in life expectancy as increasing wealth inequality while however declines in population growth decrease it. Differently from most of this literature, this paper does not look at inter-generational wealth inequality from a cross-sectional perspective (how wealth is distributed among households of different ages in the same year), but rather examines inequality between generations in terms of cohorts, i.e. how agents born in different years compare to one another at the same age.

I focus on the different economic environment faced by different cohorts, and analyse wealth inequality between generations under this light. In [Hood and Joyce \(2013\)](#), the authors show how in spite of higher real incomes at comparable ages, younger cohorts in the UK are so far not

keeping up with previous generations in building up wealth or accessing housing. The principal contribution of this paper is to reach similar findings within a cross-country empirical framework, and additionally linking demographics as a contributing factor to these developments; I also show plausible projections for wealth accumulation for a large number of countries up to the end of the century. In addition, the results of this paper provide a specific age-bracket dimension to the well-documented negative effect of ageing demographics on rates of return and positive wage effect.

1.3. DATA

1.3.1. Sources

The analysis performed in this paper makes use of a large panel dataset. Data on macroeconomic aggregates such as GDP and other national accounts components, the capital stock, employment etc. are taken from version 9.1 of the Penn World Tables (PWT 9.1) ([Feenstra et al., 2015](#)). The Penn World Tables are also the source of data on productivity, capital share of income and depreciation, among others. PWT variables are used to compute the return on capital, appreciation of capital, and its aggregate rate of return as detailed below in the following subsection. The dataset includes 182 countries between 1950 and 2017.

Data on countries' population are gathered from the United Nations' World Population Prospects 2019 ([United Nations Department of Economic and Social Affairs, Population Division, 2019](#)). These include projections on the population of countries at each year from 1950 up to 2100 and are broken down by five-year age groups. This allows for a detailed analysis of the evolution of the demographic structure of a country over time, both in terms of growth and of age distribution.

I combine the PWT 9.1 and UN population prospects data to create a panel dataset of 64 countries, spanning between the years 1950 and 2100, where the Penn World Tables data stops at year 2017. The countries were chosen based on having a sufficient number of time series observations for the most relevant variables to the analysis conducted in this paper; having typically better data, most of the advanced economies are represented in this final sample, with however around 20 developing economies (depending on definitions) represented as well. The final panel used for regression analysis

in Section 1.4 contains about 3,000 observations, or an average of almost 47 years for each of the 64 countries in the panel.

The economies included in the panel are crucially at very different stages of aging and population growth, and have quite different population distributions, life expectancies and birth rates. For the countries with the most complete data, this can go back as far as 1950 and up to 2017, allowing to cover a time span of almost seventy years, during which many economies, especially in the developed world, experienced radical shifts in their population dynamics, from growth, birth rates, age distribution and migration patterns. During these years, life expectancy also quickly increased all across the world.

Before proceeding to the empirical analysis of these data, the next two subsections explore a few data issues more in detail.

1.3.2. Total Return on Capital

This paper investigates how demographics affect wealth inequality through the lens of rates of return. Computing an exact measure of real return on wealth requires time series of price indices as well as yields (i.e. dividends, interest, rents, etc.) for all asset classes, and the relative weight of each asset class in a typical household's portfolio. This is not generally possible apart from a handful of countries, and some strong assumptions.

The alternative approach used here to estimate return on wealth draws from [David et al. \(2014\)](#) and makes use of macroeconomic data. I compute returns on an economy's capital by considering what would be the before-tax annual gains, realised and unrealised, accruing to a representative agent investing into the capital of such economy, as a fraction of the amount of capital initially invested. Here I consider the returns accruing to a domestic investor who invests into the "average" unit of domestic capital.¹

While wealth can take the form of various types of financial and non-financial assets, aggregating returns across asset types allows me to treat aggregate wealth, and its return, as equivalent to a generic capital good. In this way, developments in the return to capital can be equated to returns to the portfolio of a household with the average wealth allocation. Ultimately, returns on each type of asset can be associated with returns generated with the use of physical capital somewhere, with housing assets being the most

¹This also removes complications related to changes in relative prices of consumption goods between the investor's country and the country where capital is invested.

straightforward. The return on government bonds, for instance, is linked to the domestic capital held by the government, as well as the collection of taxes that are again collected from income generated from the use of private domestic capital.

Treating domestic agents' wealth and an economy's capital as equivalent crucially relies on the assumption of closed capital markets, where all the wealth of domestic agents is invested into domestic capital, and there is no capital flow. Any differential between the return on domestic assets, and the return on the wealth of domestic agents, is due to the openness of capital and financial markets, and the extent to which agents own higher- or lower-yielding foreign capital in their portfolios.² While capital flows may substantially affect this differential, there is very strong evidence of home bias in equity markets (see for instance [Cooper et al. \(2013\)](#))³ and in government bond markets, while the largest component of wealth in most households' portfolios (and largest component of any economy's capital stock), i.e. residential housing, is in very large part held by domestic agents. Indeed, the wealth owned by the median household in most countries tends to be even more home-biased than the average, as typical households have very little exposure to equity markets and hold most of their wealth in domestic housing. For the above reasons, the rest of this paper treats returns on an economy's capital and on its citizens' wealth as essentially equivalent.⁴

A measure of total return on domestic capital can be computed as the sum of two components:

- Income flows accruing to capital, such as payment of interest, dividends, or rents, depending on the type of asset, but also profits retained by corporations. These are equivalent to total income accruing to capital, divided by the stock of capital, and can be retrieved from the data as $Y_t^K = \alpha Y_t / K_t$, with α being the capital share of income. This can be interpreted as the marginal product of capital.
- Real changes in the price of capital, relative to consumption goods

²Recall the findings in [Krueger and Ludwig \(2007\)](#)'s model.

³Compare also [Mishra \(2015\)](#), who finds an average equity home bias of 0.77 (with 1 being complete home bias) across 42 countries.

⁴Even when domestic agents invest in other economies, they tend to be biased towards countries with similar characteristics to their own, including in terms of demographic dynamics. To the extent to which demographic factors influence rates of return, as this paper argues, investing in a country with similar demographics is almost equivalent to investing domestically, further supporting this argument.

(appreciation of capital), $(P_{t+1}^K - P_t^K) / P_t^K$, where P_t^K represents the real price of capital goods relative to consumption goods, at time t .

Total return on capital at time t (TRK_t) is thus:

$$TRK_t = \alpha \frac{Y_t}{K_t} + \frac{P_{t+1}^K - P_t^K}{P_t^K} \quad (1.1)$$

with all the necessary data available within the Penn World Tables.


While it is informative to think of returns on capital in terms of these two components, it turns out that what matters for the questions posed by this paper is exclusively the first component. While Y^K in the long term depends on productive capacity considerations alone, which are slow moving, persistent and closely interconnected with demographic developments, the appreciation of capital appears to be a short-term phenomenon more related to the business cycle.

Appreciation of capital averages at almost exactly 0% per year for the countries in the panel, signifying that there is no general evidence of capital goods increasing in value relative to consumption goods. However, it is interesting to note that while capital has not appreciated significantly across the world in the last 70 years, this net-zero appreciation hides the fact that prices of one component of it, residential and non-residential structures, have been steadily rising while the other types of capital (machinery, transport equipment, and other) have become relatively cheaper over the years.

To illustrate this, Figure 1.2 shows how the average price level of the aggregate capital stock in real terms has broadly remained unchanged, while there has been a clear trend upwards in the prices of buildings, and downwards in the other types of capital. Although the following section will be devoted to the analysis of developments in returns to aggregate capital only, this fact has potentially very important consequences for developments in wealth inequalities and should be kept in mind. The potential effects of such an appreciation of the value of buildings (and especially housing, which represents the largest component of most households' portfolios) on wealth inequality are discussed briefly later, however further research is warranted to fully evaluate the extent to which demographic developments have anything to do with these trends.⁵

⁵In Lisack et al. (2017), population aging does lead to an increase in house prices.

Figure 1.2: *Average Real Capital Price Index, and Components*



Real Capital Price Index (RKPI), for aggregate capital, machinery and equipment, transport equipment, and residential and non-residential structures (other assets omitted); unweighted average values across the countries in the panel (2011 = 1). For each country, the indices are calculated as the capital stock price index (PWT, 2011 = 1) divided by the GDP deflator (2011 = 1).

1.3.3. Population and Demographic Distribution

Populations change extremely slowly, which makes estimating the effects of such changes on economic variables, which typically have considerably more volatility, particularly challenging. Demographic changes are felt and experienced only over multiple years and even decades, as opposed to months and quarters. Nevertheless, demography remains at all times one of the most fundamental factors underpinning the performance and direction of an economy.

For this reason, while births, deaths and migration are what ultimately drive populations, because their effect is only seen years later, this paper mainly looks at demographics through the lens of the age-structure of an economy's population, or the population pyramid. Population pyramids are the result of a combination of past birth rates and death rates, evolutions in life expectancy and migration, but adequately represent where an economy's population is at the moment and where it is going. I will make use of these for most of my analysis below.

The data from the UN Population Prospects include countries' populations by five-year age groups, by number of persons and as a share of the total population in the country. This results in 18 5-year age groups (85+ year-olds are here aggregated into one only group due to the small size of the population in this age bracket), starting from 0 to 4 year-olds, and so on.

I take the working-age population to be the population between the ages of 20 and 64, inclusive. This appears to be the most appropriate choice when taking account of cross-country differences (and even differences within a country throughout the period under analysis) regarding pension age and typical age at which people get their first job.

In regression analysis presented below I seek to determine the effect of the relative size of each age-group population share on a number of variables. To do so, the first empirical approach could be to add each of these shares into a regression. However, population shares are collinear, and their sum is by definition 1 at all times. Removing one of these shares from the regression could potentially solve the issue with collinearity, however population shares that are close to one another are typically highly correlated due to the slow-moving nature of population aggregates, leading to unstable parameters. In addition, all results would have to be interpreted relatively to the omitted group rather than on their own,

which adds extra complications, and using 17 variables in a regression significantly lowers the number of degrees of freedom.


In order to avoid such issues, I construct a different set of variables employing an alternative method for representing the age distribution, i.e. using a way of parameterising the shape of the population pyramid by fitting a polynomial curve to the population shares by age group. A similar method was first introduced in [Fair and Dominguez \(1991\)](#), and as used in [Kopecky \(2022\)](#) it solves both collinearity and correlation issues mentioned above.⁶ In an efficient way of approximating the whole age distribution with a cubic function, the 18 age group population shares can be converted into three demographic variables, D_1 , D_2 and D_3 , that retain information on the shape of the population pyramid. The values of D_1 , D_2 and D_3 for country i and year t are computed as follows from the 18 age-group population shares:

$$\begin{aligned} D_{1,i,t} &= \left(\sum_{j=1}^{18} j \frac{N_{j,i,t}}{N_{i,t}} \right) - \frac{\sum_{j=1}^{18} j}{18} \\ D_{2,i,t} &= \left(\sum_{j=1}^{18} j^2 \frac{N_{j,i,t}}{N_{i,t}} \right) - \frac{\sum_{j=1}^{18} j^2}{18} \\ D_{3,i,t} &= \left(\sum_{j=1}^{18} j^3 \frac{N_{j,i,t}}{N_{i,t}} \right) - \frac{\sum_{j=1}^{18} j^3}{18} \end{aligned} \quad (1.2)$$

where $N_{j,i,t}$ is the population belonging to age group j in country i and year t and $N_{i,t}$ is total population. These three variables can be used in regression analysis, and the implied coefficients of each age-group population share $N_{j,i,t}/N_{i,t}$ on the dependent variable can be subsequently recovered from the coefficients on D_1 , D_2 and D_3 as further detailed in Section 1.4.2. Figure 1.3 shows an example of how the polynomial representation of the demographic distribution based on D_1 , D_2 and D_3 fits the actual population pyramid. Because the fitted polynomial is of third order, at times it fails to capture double humps that would be better approximated by a higher-order polynomial (like in the example), however the fit is very close to the original distribution and the correlation between the original population shares and the implied shares based on the three demographic variables is

⁶See [Kopecky \(2022\)](#) for a detailed explanation of the rationale and derivation of these variables.

Figure 1.3: Example of a demographic distribution and its polynomial representation based on the D_1 , D_2 and D_3 variables. USA, 2021 (as from UN World population prospects 2019).



very high.⁷

1.4. DEMOGRAPHICS AND RATES OF RETURN

1.4.1. Working-age Population and Productive Capacity

Economic theory and abundant literature suggest that, as the share of population in working age falls, returns on productive capital should follow the same path, as the productive capacity of the economy, per unit of capital, decreases as a result of the decline in labour input. This fact can be shown empirically with a simple exercise. Start from the simple assumptions that the rate of return on capital r^k is a negative function of the capital stock per hour worked K/H and a positive function of productivity

⁷The average correlation between the population shares in the original population pyramids and those implied by the D_1 , D_2 and D_3 variables is 94.04%, ranging from 82.80% for the 40 to 44 age-group to 99.23% for the 5 to 9 age-group. I have tested and confirm that using a fourth-level polynomial (i.e. adding a D_4 variable) does not qualitatively change any of the results shown below in any meaningful way.

A and the capital share of income α .⁸

$$\begin{aligned} r^k &= f\left(\frac{K}{H}, A, \alpha, \dots\right), \\ \frac{\partial r^k}{\partial(K/H)} &< 0, \\ \frac{\partial r^k}{\partial A}, \frac{\partial r^k}{\partial \alpha} &> 0 \end{aligned} \quad (1.3)$$

To show the effect of demographic factors on r^k , K/H can be broken down into: $K/N \times L/H \times N/L$, where N is total population, and L is the number of persons employed. In this expression therefore, each factor represents respectively (i) the capital stock per capita; (ii) the inverse of average annual hours worked per person employed; and (iii) the inverse of persons employed as a share of the population. Substituting this into the above equation results in:

$$r^k = f\left(\frac{(-)}{K}, \frac{(-)}{N}, \frac{(-)}{L}, \frac{(+)}{H}, \frac{(+)}{A}, \frac{(+)}{\alpha}, \dots\right) = f\left(\frac{(-)}{K}, \frac{(+)}{H}, \frac{(+)}{L}, \frac{(+)}{N}, \frac{(+)}{A}, \frac{(+)}{\alpha}, \dots\right) \quad (1.4)$$

This shows, unsurprisingly, how we should expect a higher share of population at work (L/N) to increase return on capital, and a higher number of hours worked per person employed to also lead to higher returns. L/N can be further decomposed into $L/N_w \times N_w/N$, where N_w is the working-age population. In other words, this corresponds to: (i) persons employed as a share of the working-age population (or the employment rate) times (ii) the share of the working-age population in total population. This finally gives us an expression that directly expresses the return on capital in terms of the working-age population share, where there is a positive relationship between r^k and N_w/N :

⁸Intuitively, working a unit of capital for longer (= lower K/H) is going to generate more income, and the same holds if productivity is higher or if a larger share of income is distributed to capital. These conditions hold for instance if we assume a standard Cobb-Douglas production function (with the labour input being hours worked, H) of the form $Y = AK^\alpha H^{1-\alpha}$, where we would have $MPK = \alpha \frac{Y}{K} = \alpha A \left(\frac{K}{H}\right)^{\alpha-1}$. However, the above formulation allows for a wider range of potential production functions.

$$r^k = f \left(\frac{(-)}{K}, \frac{(+)}{H}, \frac{(+)}{L}, \frac{(+)}{N_w}, \frac{(+)}{N}, \frac{(+)}{A}, \frac{(+)}{\alpha}, \dots \right) \quad (1.5)$$

This very simply highlights how, *ceteris paribus*, a higher proportion of the population in their working age should increase the rate of return on capital (all else equal) as it will feed into higher labour supply (hence lower K/H), increasing income.

I want now to test these facts empirically, as all the relevant variables in 1.4 and 1.5 can be easily obtained from the data. The above derivation suggests that a regression model of the following form might be appropriate within a fixed-effect panel framework to test the validity of equation 1.4:

$$\begin{aligned} \ln(r_{it}^k) = & \beta_0 + \beta_1 \ln(\alpha_{it}) + \beta_2 \ln(A_{it}) + \beta_3 \ln(K_{it}/N_{it}) + \beta_4 \ln(H_{it}/L_{it}) + \\ & + \beta_5 \ln(L_{it}/N_{it}) + \gamma_i + \delta_t + u_{it} \end{aligned} \quad (1.6)$$

where i and t represent country and year, γ_i are country fixed-effects and δ_t are year fixed-effects, and all variables are expressed in logarithms. A corresponding equation can be derived for 1.5. I test 1.6 using the panel dataset introduced in Section 3.2 and display the results in Table 1.1. The estimation employs robust standard errors clustered at country level, and year fixed-effects in specifications (2) and (4). The regression results show that what expressions 1.4 and 1.5 suggest in terms of the signs of the effects of each component of K/H , as well as A and α , on r^k holds empirically. As displayed in Table 1.1, all the assumptions arising from the simple exercise above are confirmed with significance, whether or not year fixed-effects are added, with the coefficients showing the expected signs and being considerably large. Of particular interest for this paper, a greater proportion of the population in their working age appears to have a large and positive effect on return on capital, with a 1% increase in the working-age population to population ratio increasing the return on capital by 0.49% (0.57%).

Although the results in Table 1.1 are encouraging as they match the equation formulation of the problem, there might be potential issues with such specification. The variables used for the regression in Table 1.1 are

Table 1.1: Fixed-effects panel estimation of return on capital based on equations 1.4 and 1.5.

	(1) $\ln(r^k)$ β (s.e.)	(2) $\ln(r^k)$ β (s.e.)	(3) $\ln(r^k)$ β (s.e.)	(4) $\ln(r^k)$ β (s.e.)
$\ln(\alpha)$	0.909*** (0.09)	0.930*** (0.09)	0.909*** (0.09)	0.932*** (0.09)
$\ln(TFP)$	0.673*** (0.09)	0.614*** (0.10)	0.674*** (0.09)	0.617*** (0.10)
$\ln(K/N)$	-0.352*** (0.04)	-0.322*** (0.05)	-0.354*** (0.04)	-0.333*** (0.05)
$\ln(H/L)$	1.061*** (0.23)	0.902*** (0.26)	1.058*** (0.23)	0.879*** (0.26)
$\ln(L/N)$	0.478*** (0.11)	0.479*** (0.13)		
$\ln(L/N_w)$			0.471*** (0.14)	0.442*** (0.14)
$\ln(N_w/N)$			0.492*** (0.16)	0.570*** (0.19)
Constant	-4.804** (1.98)	-3.878* (2.09)	-4.753** (1.99)	-3.540 (2.13)
Country fixed-effects	✓	✓	✓	✓
Year fixed-effects		✓		✓
R^2	0.632	0.656	0.632	0.657
deg. fr.	63	63	63	63
# obs.	3000	3000	3000	3000

slow-moving, some of them are trending upwards (e.g. K/N , TFP) or downwards (H/L) in most countries, and they might be cointegrated. A visual inspection of the data and the use of unit root tests for the above variables at the country level seem to confirm that for most of the countries, the variables included in the above specifications are generally non-stationary, and in most of cases they have a trend. (Panel unit root tests are more contradictory in this regard however, with some types of tests rejecting the null hypothesis that the variables are non-stationary in all panels for some of the variables.) There is a strong possibility that the variables might therefore be cointegrated, and indeed in the case of a few countries, a Vector Error-Correction model with one cointegrating equation would well fit the data.

To address these issues I have separately estimated specifications (1) and (3) of Table 1.1 (i) with all variables differenced once, and (ii) as an Arellano-Bond dynamic panel. These give similar results: in both regression models, the coefficients on all the explanatory variables remain highly significant, large and with the expected signs. The results from Table 1.1 do not appear to be the consequence of a spurious relationship, and can be reasonably taken as valid.⁹

Another concern lies in the choice of variable for productivity. To generate a measure of Total Factor Productivity, as calculated in PWT and used for Table 1.1, a certain underlying production function in the economy needs to be assumed; this production function might not be the true underlying function and bias the results through the addition of TFP in the regression. At the same time, omitting TFP from the regression would clearly lead to omitted variable bias. In the absence of any specific reason for choosing a different productivity measure (which would still be reliant on some different assumption), and because of its widespread and established use in the literature, in the analysis that follows I will continue using TFP as produced by the Penn World Tables. However, in Appendix A, I calculate alternative productivity measures and repeat the exercise done for Table 1.1 (as well as the following Section). The results suggest that using TFP remains appropriate.

Finally, while I argue above that what matters most for domestic capital accumulation is the return of domestic capital (due to high home bias in investments, especially for the median household in most countries),

⁹Results are available but are omitted for brevity.

nevertheless it needs to be recalled that free capital mobility, where present, allows agents to look for investment opportunities abroad and own foreign capital. Such capital will have a different rate of return than domestic capital, in part also due to differences in demographics. As a result, to the extent to which they invest abroad, domestic agents can be insulated from the domestic rate of return, and partly 'import' different rates of return from abroad. As such, the relationship between domestic demographic developments and the rate of return on capital actually enjoyed by domestic agents, as found here and in the analysis below probably represents an upper bound, especially in economies characterised by high capital mobility.

1.4.2. Age Distribution and Capital Accumulation

While the above analysis suggests that the working-age population share alone has a direct impact on returns due to productive capacity considerations, economic theory also suggests that the *structure* of the population, i.e. the shape of the population pyramid, is in turn likely to affect some of the other factors into which the return on capital was decomposed in equations 1.4 and 1.5. Ignoring these channels can potentially significantly over- or underestimate the effect of the demographic distribution on returns. Indeed, even among people of working age, agents of different ages can have quite different effects on the economy.

Every item in equation 1.4 might be dependent on demographics. In particular, K/N stands out in this sense. An abundant literature on the consumption, savings, and wealth accumulation patterns of households throughout their life cycle shows how households typically start their working life with limited assets (or in debt); they gradually build up their wealth through the years by accumulating savings from their incomes, which also typically rise through the working years, reaching a peak after their fifties; finally, once they retire, they slowly run down their assets. Such a pattern of behaviour has been frequently observed empirically.

While every person's wealth accumulation journey is different and mostly independent, this pattern of behaviour aggregates across the population such that if an economy's population, all else equal, has relatively more people of an age at which life-cycle wealth is at its peak, then the overall level of capital per person in the economy will be higher. Note how, as K/N came with a negative coefficient on returns in Table 1.1, that implies that the

return on capital will be lower in such an economy due to the abundance of capital relative to the population, which is in turn due to the simultaneous large holdings of wealth by a large cohort of individuals.

In addition to capital per person, the demographic distribution is likely to have a major impact on the employment rate of the population (L/N) as well. Different age groups have vastly different labour force participation rates, with some groups not participating in the labour force at all, while even within the working-age population, individuals of different ages have different likelihoods of being in employment.

Finally, although Total Factor Productivity, or its growth, might also be higher *ceteris paribus* with a relatively younger workforce, as has been argued before in the literature (see Section 1.2), such a relationship with the age distribution is less clear and established, and therefore TFP is taken as exogenous to the demographic distribution in the analysis that follows.

To obtain an estimate of the combined effect of the age distribution on the rate of return, while accounting for the wealth accumulation and employment rate channels (i.e. the effects of the age distribution on K/N and L/N), I estimate a panel instrumental variable regression based on specification (2) of Table 1.1, where however $\ln(K/N)$ and $\ln(L/N)$ are instrumented by the demographic variables D_1 , D_2 and D_3 (calculated as explained in section 1.3.3), plus year fixed-effects.

The shape of the population pyramid is an ideal instrument, as an economy's population structure is almost entirely independent of *current* economic conditions, due to the very slow-moving nature of population dynamics. Even migration and major events that affect both the economy and demographics (pandemics such as COVID-19, for instance), have a negligible impact on the population pyramid in the year in which they occur. However, it is not unreasonable to think that there might be a relationship between, for instance, the employment rate and hours worked, and the number of children a couple decides to have (think of time and monetary commitments), or that migration flows (both inward and outward) are largely dependent on the state of the economy. These and other developments might not affect the population distribution immediately, but they may do so over longer horizons. As such, there might be an endogenous component to the demographic variables; it is however not possible to fully disentangle it.

Table 1.2 displays the results of the first-stage regressions on $\ln(K/N)$ and $\ln(L/N)$.

Table 1.2: Instrumental-variables estimation of r^k . First-stage regressions.

	(1) $\ln(K/N)$ b/se	(2) $\ln(L/N)$ b/se
D_1	0.260 (0.21)	0.731*** (0.09)
D_2	0.097*** (0.03)	-0.046*** (0.01)
D_3	-0.006*** (0.00)	0.000 (0.00)
$\ln(TFP)$	0.800*** (0.02)	-0.118*** (0.01)
$\ln(\alpha)$	-0.099*** (0.03)	0.036*** (0.01)
$\ln(H/L)$	0.924*** (0.08)	-0.104*** (0.03)
Constant	3.851*** (0.58)	-0.309*** (0.24)
Country fixed-effects	✓	✓
Year fixed-effects	✓	✓
R^2 (overall)	0.556	0.423
# obs.	3000	3000


The demographic variables D_1 , D_2 and D_3 are jointly significant at the 1% level for both $\ln(K/N)$ and $\ln(L/N)$, indicating that the shape of the population pyramid that they represent has itself an impact on these variables. They also jointly explain a large proportion of the variation in the variables they are instrumenting: removing all three variables reduces the overall R^2 from 0.556 to 0.048 in the case of the $\ln(K/N)$ first-stage equation, and from 0.423 to 0.032 in the case of the $\ln(L/N)$ equation. Removing them one at a time instead leads only to negligible changes to the overall R^2 : this is because, in any combination, any two of D_1 , D_2 and D_3 still approximates the age distribution well. This shows again that the three variables only have a meaningful interpretation jointly, and that adding a D_4 for a four-degree polynomial approximation of the population distribution is unnecessary; this fact holds for this as well as all specifications below that contain the three D variables.

The implied coefficients of each age bracket on $\ln(K/N)$ and $\ln(L/N)$ can be recovered from the coefficients on D_1 , D_2 and D_3 , based on how these were calculated from the relative sizes of the age brackets (see Section

1.3.3). Recalling that there are 18 age brackets in the data (17 5-year age groups, plus a group including all the population above 85 years old), the implied coefficient $\beta^{j, \ln(K/N)}$ of age group j on $\ln(K/N)$, given the coefficients $\beta^{D_1, \ln(K/N)}$, $\beta^{D_2, \ln(K/N)}$ and $\beta^{D_3, \ln(K/N)}$ from the first-stage of the instrumental variables regression in Table 1.2, can be calculated as:

$$\beta^{j, \ln(K/N)} = -\left(\frac{\beta^{D_1, \ln(K/N)}}{18} \sum_{k=1}^{18} k + \frac{\beta^{D_2, \ln(K/N)}}{18} \sum_{k=1}^{18} k^2 + \frac{\beta^{D_3, \ln(K/N)}}{18} \sum_{k=1}^{18} k^3\right) + j\beta^{D_1, \ln(K/N)} + j^2\beta^{D_2, \ln(K/N)} + j^3\beta^{D_3, \ln(K/N)} \quad (1.7)$$

Figure 1.4: *Effect of each age group on $\ln(K/N)$ and $\ln(L/N)$*



with $j = 1$ representing the age group of persons aged 0-4 years old, $j = 2$ representing 5 to 9 year-olds, and so on at 5-year increments. The equivalent can be done for $\ln(L/N)$. Figure 1.4 displays graphically the backed-out coefficients from the first-stage regressions for each age group population share, derived from the coefficients on D_1 , D_2 and D_3 as for the equation above. The figure shows how, in accordance with the life-cycle hypothesis, agents in their late working-age and early retirement

increase capital per person in the economy, as these are the years in which the highest level of wealth is accumulated before asset positions begin to be sold as individuals age further; conversely, a larger proportion of children and the elderly in the population tends to decrease capital per capita. Unsurprisingly, the coefficients on the employment rate are instead positive for all age groups of working age. The peak positive impact of an age-group on employment comes earlier than for capital accumulation: while 55 to 64 year-olds have the largest impact on $\ln(K/N)$, 35 to 44 year-olds have the largest impact on employment.

Table 1.3: Instrumental-variables estimation of r^k . Second-stage regression.


	(1) $\ln(r^k)$ b/se
$\ln(K/N)$	-0.443*** (0.02)
$\ln(L/N)$	0.854*** (0.06)
$\ln(TFP)$	0.773*** (0.04)
$\ln(\alpha)$	0.901*** (0.02)
$\ln(H/L)$	1.015*** (0.06)
Constant	-3.148* (0.44)
Country fixed-effects	✓
Year fixed-effects	✓
# obs.	3000

The second-stage of the IV panel regression (Table 1.3), with $\ln(K/N)$ and $\ln(L/N)$ instrumented by D_1 , D_2 , D_3 and year fixed-effects, gives results that are very similar to those from specifications (1) and (2) in Table 1.1. However, this specification allows to produce an estimate of the effects of the whole age distribution on $\ln(r^k)$. By multiplying the implied coefficients of the age groups on $\ln(K/N)$ and $\ln(L/N)$ from the first-stage regressions as displayed in Figure 1.4, by the coefficients of $\ln(K/N)$ and $\ln(L/N)$ themselves on the rate of return from the second-stage regressions, we can estimate the effect of each age group on the return on capital through wealth accumulation channel, and the employment channel. The sum of the two (see equation 1.8) provides an estimate of the overall impact of each age group on rates of return through the combined effect of these

two channels. The results are displayed graphically in Figure 1.5.

$$\beta^{j, \ln(r^k)} = \beta^{j, \ln(K/N)} \beta^{\ln(K/N), \ln(r^k)} + \beta^{j, \ln(L/N)} \beta^{\ln(L/N), \ln(r^k)} \quad (1.8)$$

Figure 1.5: Combined effect of each age group on rates of return via $\ln(K/N)$ and $\ln(L/N)$



These results are very intuitive. Not participating in the labour force, the effect of children and retired persons on returns via the labour force participation channel is negative. At the same time, children and young adults, and persons in late retirement contribute a positive effect on returns via capital accumulation due to their low levels of accumulated wealth, for opposite reasons. This results in an overall positive effect on returns by young agents overall, with the peak positive impact at about 20 years of age and then steadily decreasing.

As agents enter the labour force, they increase returns to capital through their labour, but while their savings gradually accumulate, steadily increasing K/N , the negative impact on returns from wealth accumulation rises and finally overcomes the positive effects of labour force participation in the 45-49 age bracket, after which the combined effect becomes negative. With the labour-force participation effect now falling steadily and wealth accumulation remaining very high, the 70-79 age groups become the ones

that most adversely affect returns. While agents above 80's typical asset decumulation leads to a positive effect on returns through the saving channel, this is not sufficient to fully compensate for the negative effect through the employment channel.

The median ages of advanced economies are now moving towards the peak earning years of the late working age (> 50), while the shares of people in their teens, twenties and thirties are rapidly falling due to falling birth rates, and the share of over-65s rises steadily. This results in lower returns on capital due to the combined effects of (i) a falling working-age population overall, and especially the younger end; and (ii) a higher proportion of people with substantial wealth, relatively less productive and close to retirement bringing up K/N .

Based on this analysis and projected demographic developments across the world, advanced economies are set to continue to experience lower and falling rates of return for many years, while an increasing number of developing economies which are starting to experience a demographic transition will also see their rates of return fall in the next decades.

These factors highlight a stark generational imbalance between those currently in or close to retirement, who were able to build up considerable amounts of wealth in the past, when the younger structure of the population, high growth and relative scarcity of capital allowed for high returns; and a cohort of young individuals facing a future of low returns impairing their ability to build wealth themselves. The effects of falling returns on inter-generational wealth imbalances are explored further in the next Section.

1.4.3. Returns to Wealth and Inequality Between Generations

The previous Section showed how a relatively young population positively affects rates of return on capital, as in young economies labour is abundant and capital relatively scarce. Similar results are well-established in the literature on secular stagnation, in particular with respect to the effects of population ageing on long-term equilibrium interest rates and risk-free rates.

In an ageing economy, as rates of return gradually decline, including due to demographic factors, and wealth becomes increasingly concentrated

in the hands of older age groups, the accumulation of wealth becomes more and more difficult for younger generations of workers. At the same time, the lower production of income out of wealth due to declining rates puts an increasing number of retirees with limited resources into financial difficulties in the absence of generous, but increasingly expensive public pension schemes. In such an economy, production by the current, shrinking young labour force needs to generate enough returns and tax revenues to sustain the incomes of a rising number of retirees, through public pension schemes as well as capital gains.

To illustrate the link between population dynamics and wealth inequalities between generations, the results from Section 1.4.2 can be used to compute the overall effect of the aggregate demographic distribution of an economy on r^k , and therefore on total rates of return in the long term. To do so, for every country and year, the implied coefficients of each age bracket on $\ln(r^k)$, as illustrated graphically in Figure 1.5, are multiplied by the share of each age bracket in the total population, to give the impact of such age bracket on $\ln(r^k)$. These are then summed up in order to recover the impact of the aggregate demographic distribution itself. The effect of the demographic distribution on $\ln(r^k)$ in year t on country i is therefore calculated as:


$$\beta_{i,t}^{DD} = \sum_{j=1}^{18} \beta^{j,\ln(r^k)} \frac{N_{j,i,t}}{N_{i,t}} \quad (1.9)$$

where $N_{j,i,t}$ is population in age bracket j in country i and year t ($N_{i,t}$ is the total population in the same country and year), and the β^j 's are calculated from the coefficients on D_1 , D_2 and D_3 as in equation 1.7.

By using the values for $\beta_{i,t}^{DD}$ across countries and years, I can estimate the differential impact of age distributions on the accumulation of wealth of different generations, and explore the extent to which the impact of age distributions on returns may represent a significant source of wealth inequality between generations, in the present and future.

Figure 1.6 displays the estimated effect of the overall demographic distribution on returns on capital for a number of advanced economies. This is computed from $\beta_{i,t}^{DD}$ as per equation 1.9; as $\beta_{i,t}^{DD}$ is the effect of the demographic distribution on the natural logarithm of r^k , the effect on r^k itself is calculated as:

Figure 1.6: *Impact of the demographic distribution on return on capital (r^k), selected advanced economies.*



$$DD_{i,t}^{effect} = r_{i,t}^k - e^{\ln(r_{i,t}^k) - \beta_{i,t}^{DD}} \quad (1.10)$$


While it turns out that such an effect has been considerably positive for most of the time series and countries under consideration, due to young and growing populations throughout most of the twentieth century, the trend is downwards everywhere (including countries not in the figure), with aging populations and lower birth rates moving the demographic distribution towards age groups that lower r^k .

Figure 1.7 further illustrates the evolution of rates of return in China, Italy, Japan and the USA,¹⁰ the estimated impact of the demographic distribution on these returns throughout the years, and finally what returns would have been in the absence of such demographic factors. As is evident for the countries displayed in Figure 1.7¹¹, rates of return after removing

¹⁰Results for the other countries are available. These countries were chosen because they include two of the oldest populations in the world (Italy and Japan), and the two largest economies (USA and China). The case of the USA is also interesting as it remains one of the youngest advanced economies (partly due to high immigration) and did not experience as dramatic a demographic shift as most of Europe; on the other hand, the population of China is aging at a fast rate also due to its former one-child policy.

¹¹But also for the rest of the countries

Figure 1.7: Returns to capital and demographic effect: CHN, ITA, JPN, USA



demographic effects (i.e., $r^k - DD_{i,t}^{effect}$) appear to be significantly more stable throughout the years in individual economies (with little to no sign of them trending up or downwards in most countries), and also more similar between economies, as countries with higher rates overall also tended to have younger populations.¹² Thus, for a large part due to demographic developments, rates of return saw a significant decline in recent decades as the positive contribution of the population structure (a “demographic dividend” that certain generations have enjoyed) fell across the world. In a few countries (e.g. Italy and Japan), the demographic distribution is now weighing on rates of return, which would have been higher otherwise, with the turning point occurring in both cases soon after the year 2000. As also shown in Figure 1.6, more economies are now in this position due to their ageing populations, or are trending towards it.

Figure 1.7 also presents a forecast of the evolution of rates of return through the end of the century. This is done by assuming that $r^k - DD_{i,t}^{effect}$, which as discussed has been relatively constant throughout the sample period, remains at a level equal to its country average from the year 2000 to the latest observation. $DD_{i,t}^{effect}$ is then calculated by obtaining D_1 , D_2 and D_3

¹²The standard deviation of r^k across the entire sample is 0.083, while for r^k with demographic effects removed it reduces to 0.052.

based on the UN population projections. r^k is then finally the sum of the two. In countries where the demographic transition is already at an advanced stage, rates of return are expected to fall further (but only slightly) for a few more decades, reaching their lowest point at around 2040/2050, to then stabilise around that level, below the rate of return that would have been achievable in the absence of any demographic contribution, and significantly below the rates of return experienced in the second half of the twentieth century. Economies at an intermediate stage of the demographic transition (e.g. China) are expected to follow the same path with a few decades' delay; their rates of return are seen declining rapidly in the next years.

Most developing economies are still expected to enjoy a high positive contribution to returns, thanks to their relatively young population, for some time to come; however, in almost all cases, rates of return are expected to decline significantly and steadily for the most part of the century. The contribution of the population of China to China's rate of return on capital, for instance, stood around 2.9% in 2017, while it is expected to reach a level of -0.7% by around 2070; this implies a decline in rates of return there of almost 7 basis points per year for over 50 years, due to a rapidly declining birth rate and rising life expectancy. Finally, countries which are expected to maintain relatively higher birth rates and/or higher levels of net migration (e.g. the US) will see rates of return declining at a much slower pace.

Young, growing and dynamic populations throughout the twentieth century have helped maintain high returns to capital/wealth, thanks to a high share of young and productive workers and a relative scarcity of capital; these high returns have allowed generations of middle-class households to accumulate sufficient wealth for a decent retirement. Indeed, throughout this time, a pyramid-shaped population distribution has sustained a pyramid scheme of sorts, where the returns on wealth of a cohort were guaranteed by income produced by the next, larger and younger cohort of workers.

As life expectancies increased, and birth rates fell, population pyramids in advanced economies started stabilising, as new generations were no longer larger than the previous ones and more people survived to older ages; in recent years, population pyramids have started to become inverted, with younger generations shrinking more and more. Recall Figure 1.1 and the developments in the population of Japan, where the demographic shift is


currently at the most advanced stage globally. In Japan and economies with similar demographic developments, what was before a “pyramid scheme” where increasingly large younger generations supported the retirement incomes and wealth of a small older population, has slowly morphed into an *inverted pyramid scheme*, where it becomes quite hard for a shrinking young population to now do the same.

Whether an agent is young, middle aged, or retired, the age of everyone else around them matters greatly to their wealth accumulation, via the effects of the demographic distribution on returns on capital as seen above. But how much more wealth could an agent who started saving in the 1960s, with a growing and young population, accumulate over the next 20 years compared to a similar agent who started investing in the 2000s, and how much of this difference can be attributed to demographics?

If returns are compounded annually, a small difference in annual returns can lead to very significant wealth accumulation paths across a working life. With the “demographic dividend” described here starting to fall across the developed world around the 1970s and bringing down rates of return with it, each subsequent generation would have been at a disadvantage compared to the previous one in terms of their wealth accumulation. Newer generations would either accumulate less wealth out of the same savings compared to their elders, or be forced to save a larger portion of their incomes to make up for the impact of the lost returns. Figure 1.8 shows how a small difference in the rates of return experienced by households in subsequent generations in the countries of Figure 1.7 lead to widely diverging wealth accumulation journeys. An Italian household who invested 1 in 1955 and reinvested the return (r^k) every year for 20 years could expect to have accumulated a real value of 30 by 1975. By 2017, an investment of 1 20 years prior would have only returned a real value of 10. Notice how this wedge was mostly due to the very favourable demographic structure just after WWII pushing up rates of return, which then fell steadily from the 70s as population started ageing. In the absence of these demographic effects, the two households would have enjoyed virtually the same cumulated returns out of their investments.

This paper largely focuses on how demographic dynamics unequally impact wealth accumulation for different generations or cohorts. However, the large imbalances between generations in the speed of wealth accumulation have an impact on wealth inequality through multiple different

Figure 1.8: Wealth accumulated over 20 years: CHN, ITA, JPN, USA



The figure shows, for each year, the wealth accumulated if a 1 unit investment was made 20 years before, and returns were each year reinvested and compounded at the prevailing rate of return. The rates of return considered are the r^k and r^k with demographic effects removed, as in Figure 1.7.

channels.¹³

The first is apparent from Figure 1.8. In general, as they more easily accumulated wealth while younger, older generations will hold a larger share of wealth. On the other hand, younger people will be further away from the average per-capita wealth level than their elders were at the same age. This leads to *inter*-generational wealth inequalities.

At the same time, falling rates of return might mitigate *intra*-generational wealth inequalities. With high rates of return, a small difference in savings between two agents of the same age can compound to become very large over decades; with lower returns, this is less of the case, leading to less wealth inequality between agents of the same age.¹⁴

¹³One of these is the decline in rates of return documented here. Greenwald et al. (2021), for instance, find that declining real interest rates lead to an increase in financial wealth inequality.

¹⁴Note, however, that regardless of initial wealth, agents would be worse off, with lower wealth (all else equal) than in the presence of higher returns. Also, this is only the case if agents invest in the same portfolio. With different portfolio compositions, agents investing in higher-return assets might benefit if risk premia increase with lower rates of return.

Lower rates of return also tend to raise asset prices. Indeed, while Figure 1.2 showed how the price of capital remained unchanged overall, this was due to a combination of falling prices of machinery, transport equipment and other forms of capital, while the price of buildings and residential structures (and with it the value of housing assets) kept increasing in real terms. With the possession of housing assets very low at young ages and low wealth percentiles, the appreciation of housing assets due to falling rates, which are in turn due to demographic developments, would increase both intra- and inter-generational wealth inequalities.

While the above is true if we consider wealth as a single aggregate asset, agents of different ages tend to invest in different asset types, such that a different demographic distribution is likely to affect the relative prices and returns of different asset types. Kopecky and Taylor (2020) suggest that aging demographics raise the equity risk premium; at the same time, housing has become more expensive in the past decades as seen in Section 1.3.2. As agents who invest in these asset types tend to be wealthier, a falling rate of return on capital (aggregate) may affect agents in the lower end of the wealth distribution more than those in the upper end, further increasing wealth inequalities.

Inter-generational transfers (bequests) further complicate the picture. An increasing proportion of wealth being held by the elderly signifies that increasing amounts of wealth will be ultimately transferred to younger generations. In a young economy, transfers might decrease inequality by redistributing wealth to a higher number of less wealthy agents. A longer life expectancy, however, means that the average age at which agents receive such a transfer increases further.¹⁵ With life expectancy at over 80 in many advanced economies, the age at which such transfers occur would tend to fall during an agent's late working age, i.e. at the peak of their wealth accumulation, further concentrating wealth in those age groups.¹⁶ Lower birth rates also mean that wealth transfers will be increasingly concentrated to a decreasing average number of heirs. All in all, in an ageing society bequests are thus likely to worsen inter-generational wealth inequality

¹⁵This effect is slightly mitigated by the fact that parents' age at a child's birth has also tended to be on the increase in ageing economies. However, life expectancies have risen by longer, meaning that, overall, the average age of receipt of bequests has risen: for instance, Eurostat data shows that the mean age of women at childbirth in Belgium has risen from 28.0 years in 1960 to 30.8 in 2020; over the same period, life expectancy increased from 69.7 to 80.8 years. Similar trends can be seen in other countries.

¹⁶Inter-generational gifts when the recipients are in their working age can alleviate these concerns, although these tend to be smaller than bequests.

(as transfers tend to go to already wealthy age groups), as well as intra-generational wealth inequality (transfers are more concentrated).

The analysis presented so far points to convincing evidence of population dynamics having a significant effect on returns, and consequently on the unequal distribution and accumulation of wealth, especially between different generations in ageing societies. The empirical evidence from this section suggests that decreasing rates of return as population ages lead to rising inter-generational inequality, if subsequent generations invest similar amounts to each other. However, does the situation change when wage levels are taken into account? The next section explores this.

1.5. DEMOGRAPHICS AND WAGES

We have seen that population ageing contributes to a decrease in rates of return. This leads to divergences in how much wealth a household can expect to accumulate out of a certain level of savings. At parity of (real) investment, current young generations are disadvantaged compared to previous generations in terms of wealth accumulation, due to lower rates of return. However, the situation may change if wages are taken into consideration: since the 1950s, real wages have steadily increased across the world, such that in most economies an average young household today can afford to invest more (in real terms) than previous generations. Net of this, are younger generations better or worse off?

A large part of increases in labour compensation in past decades can be attributed to technological advancement and higher productivity. However, the same demographic dynamics that have been discussed for returns on capital also have a bearing on the compensation of labour, where for the most part their impact on wages is opposite to their impact on returns. In particular, a smaller proportion of the population at work and a higher capital intensity both tend to lead to higher wages.

The exercise in Section 1.4.1 can be repeated for wages (per hour worked). With similar arguments as for capital, I make three assumptions about the return on labour which are consistent with 1.3 and economic theory, i.e. wages will be higher if: (i) there is more capital relative to labour K/H (with more capital, an hour of labour can produce more); (ii) productivity is higher; (iii) the capital share of income is lower. In other words:

$$\begin{aligned}
w^H &= f\left(\frac{K}{H}, A, \alpha, \dots\right), \\
\frac{\partial w^H}{\partial(K/H)}, \frac{\partial w^H}{\partial A} &> 0, \\
\frac{\partial w^H}{\partial \alpha} &< 0
\end{aligned} \tag{1.11}$$

Decomposing K/H as in 1.4 gives:

$$w^H = f\left(\frac{K}{N}, \frac{H}{L}, \frac{L}{N} \left(= \frac{L}{N_w} \frac{N_w}{N} \right), A, \alpha, \dots\right) \tag{1.12}$$

i.e., contrary to the case for rates of return on capital, we expect K/N to increase hourly wages and a higher proportion of the population at work to decrease them. The average total wage earned by a worker, w^L , would be equal to $H/L \times w^H$. Assuming that the elasticity of hourly wages to hours worked per person is less than one (workers would not be working as many hours if that was not the case), we can expect H/L to have a positive effect on w^L such that:

$$w^L = f\left(\frac{K}{N}, \frac{H}{L}, \frac{L}{N} \left(= \frac{L}{N_w} \frac{N_w}{N} \right), A, \alpha, \dots\right) \tag{1.13}$$

How do demographics affect wages through the channels in equation 1.13? We know that ageing demographics tend to increase K/N and lower L/N (recall Figure 1.4). Running the same instrumental variables regression as in section 1.4.2 (replacing r^k with w^L as dependent variable) unsurprisingly gives results that are consistent with equation 1.13, as shown in Table 1.4. The first-stage regression results are still the same as for Table 1.2 (the variables entering the first-stage regressions are the same), giving the age-profile of effects on K/N and L/N as per Figure 1.4.

As done above for Figure 1.5, multiplying these implied coefficients of each age group on $\ln(K/N)$ and $\ln(L/N)$ by the coefficients of $\ln(K/N)$ and $\ln(L/N)$ from the second-stage regression, I can calculate the impact of each age group on labour compensation ($\ln(w^L)$) through wealth accumulation, through the labour force, and the overall effect. Figure 1.9 shows this

Table 1.4: Instrumental-variables estimation of real (log) w^L . Second-stage regression.


	(1) $\ln(w^L)$ b/se
$\ln(K/N)$	0.566*** (0.01)
$\ln(L/N)$	-0.229*** (0.04)
$\ln(TFP)$	0.792*** (0.02)
$\ln(\alpha)$	-0.800*** (0.01)
$\ln(H/L)$	0.702*** (0.03)
Constant	-2.552*** (0.22)
Country fixed-effects	✓
Year fixed-effects	✓
# obs.	3000

graphically. The strongest impact is through wealth accumulation, where age groups with high accumulated wealth lead to an abundance of capital in the economy, which generates an increase in the compensation of labour by making it more productive.

A higher employment rate of the population has the opposite effect, as less scarcity of labour pushes down its cost; this channel seems however to be smaller than the wealth accumulation channel through demographics.¹⁷ Overall, all age groups from 35-39 to 80-84 lead to higher wages, all other things constant, while over-85s and under-35s reduce them. As a result, ageing economies tend to have higher wages due to demographic factors, and in particular due to the effect of the wealth accumulated by those from middle age to early retirement. Comparing this with Figure 1.5, which shows the effect of each age group on return on capital, we can see that for most age groups, having a positive impact on return on capital

¹⁷Note how in Figure 1.4, the effects of the age distribution on $\ln(L/N)$ are flatter than those on $\ln(K/N)$. This is partly because, in ageing economies, the labour force participation of the working-age population has been observed to increase (at least in part due to demographic developments) as the working-age population share itself shrinks; at the same time, labour force participation in groups outside of the typical working-age population (over 65s) has also tended to rise in these economies, also because of improving health. The overall effect of demographics on the employment rate, while certainly strong, is therefore not as stark as might seem.

Figure 1.9: Combined effect of each age group on compensation of labour via $\ln(K/N)$ and $\ln(L/N)$



means having a negative impact on compensation of labour, and vice-versa. However, crucially, the magnitudes are not quite the same for all age groups (for instance, while the 65-69 age group has a large negative impact on return on capital and a large positive impact on compensation, the 45-49 group has a high positive impact on wages but only a negligible negative impact on returns); additionally, persons aged 35-44 have a positive effect on both returns and wages, while the opposite is true for over-85s.


The overall effect of the population pyramid on the compensation of labour is shown in Figure 1.10. More specifically, the overall coefficient of the demographic distribution on $\ln(w^L)$, $\beta_{i,t}^{DD}$, is calculated as in equation 1.9. From this (due to the fact that the regression is in log form), the effect of demographics on actual real wages (in constant 2011 US dollars) can be calculated as:

$$DD_{i,t}^{effect} = w_{i,t}^L - e^{\ln(w_{i,t}^L) - \beta_{i,t}^{DD}} \quad (1.14)$$

where $w_{i,t}^L - DD_{i,t}^{effect} = e^{\ln(w_{i,t}^L) - \beta_{i,t}^{DD}} = w_{i,t}^{L_{nodem}}$ can be seen as the level of labour compensation that would have prevailed in the absence of any

effect due to demographics. Figure 1.10 shows how big $DD_{i,t}^{effect}$ is as a proportion of this $w_{i,t}^{L, nodem}$. We can see that in the 1950s to 1980s, young population pyramids, leading to higher proportions of the population at work and lower accumulated wealth per capita,¹⁸ caused wages to be up to 30 (and more) per cent lower than would otherwise have been the case. As populations in advanced economies aged, this negative effect has progressively been reducing and has now turned positive in most of these economies; in some of them, wages would now be up to 20% lower if it were not for the scarcity of labour and abundance of capital that is due to demographic dynamics.

Figure 1.10: *Impact of the demographic distribution on compensation of labour (w^L), selected advanced economies.*




Note: Calculated as $\frac{DD_{i,t}^{effect}}{w_{i,t}^L - DD_{i,t}^{effect}}$.

According to this it appears that, in the most recent years, the impact of demographics has been a major factor behind increases in real wages in advanced economies. Productivity growth has slowed down considerably, and any potential increases in annual wages per worker due to rising productivity have been in many cases almost exactly compensated by a decrease in average hours worked. On average for these countries, real wages, once demographic effects are removed, have remained unchanged

¹⁸Note that this is after accounting for long-run trends in capital accumulation, labour force participation and productivity, which are dealt with in the regressions.

since around 2000. Figure 1.11 displays this.

Figure 1.11: *Evolution of (log) real wages per person employed (actual versus with demographic effects removed), selected advanced economies.*



Note: Values are in log constant 2011 US dollars. The countries included are the same as for Figure 1.10.

1.6. WEALTH ACCUMULATION: WINNERS AND LOSERS

The previous sections have detailed how demographic dynamics have contributed to a decline in rates of return while positively affecting wages. Also due to this, the experiences of subsequent generations in ageing economies have been (and will be) substantially different. This section gathers the results shown above for returns and wages and seeks to quantify the overall effect of demographic change on the wealth accumulation journeys of different generations in economies across the world, with a look at future decades. Are there generations that stand to lose or to gain due to these dynamics? What are the implications in terms of inter-generational wealth inequality?

In an environment of increasing wages but decreasing rates of return, subsequent generations are able to save more, but can also expect to *have* to save more in order to accumulate a certain level of wealth, because returns

are lower. To determine what effect dominates in any given year and country, I perform a couple of simple exercises that intuitively illustrate how conditions for wealth accumulation have changed through time and across country, and the role of demographics.

Similar to the exercise done for Figure 1.8, I first look at the wealth accumulated out of an investment made at a point in time, but also taking the wage level into account. Specifically, I calculate how much real wealth ($W_{i,t}^{acc}$) is accumulated by time $t + 20$ out of an investment of a fraction s_r of the average wage at time t , if this investment returns the relevant prevailing rate of return (r^k) every year, and all yields are reinvested.

$$W_{i,t}^{acc} = s_r w_{i,t}^L \prod_{j=0}^{19} (1 + r_{i,t+j}^k) \quad (1.15)$$

Compare this to a hypothetical situation where there is no effect of demographics, i.e. the wage level is $w^{L,nd}$ and the return on capital is $r^{k,nd}$. The wealth accumulated would then be:

$$W_{i,t}^{acc,nd} = s_r w_{i,t}^{L,nd} \prod_{j=0}^{19} (1 + r_{i,t+j}^{k,nd}) \quad (1.16)$$

After the end of the PWT sample, estimated demographic effects can be computed thanks to the UN population projections spanning until 2100. However, a forecast for $w^{L,nd}$ and $r^{k,nd}$ needs to be made.¹⁹ As above, I take $r^{k,nd}$ as constant through 2100 for each country given its stability in the sample. In terms of wages, I formulate three very different stylised scenarios (with very different implications), with a focus on the prospects of ageing advanced economies:²⁰

- (1) $w^{L,nd}$ remains constant. This scenario is possibly the most plausible for advanced economies as it has arguably been the case since 2000 (compare Figure 1.11). The average change in $w^{L,nd}$ since 2000 in advanced economies is 0%.²¹
- (2) It grows at a level equal to its average growth (across advanced

¹⁹ w^L and r^k will then be the forecasted $w^{L,nd}$ and $r^{k,nd}$ plus the respective demographic effects calculated from D_1 , D_2 and D_3 .

²⁰No assumption about developments so far in the future can be justified, so the approach is to try and cover as many options as possible, rather than forecast.

²¹If the vast literature on secular stagnation is to be heeded, growth is unlikely to pick up significantly in the next decades, and there are several culprits, including demographic change.

economies, unweighted) since 1980. This is equal to around one third of a percentage point per year. The average trend Total Factor Productivity growth across the countries (including developing) in the sample is less than half this, while the average TFP growth since 2000 is around 0.5%. This scenario can represent an average country, although this growth rate is slightly optimistic for advanced economies given the slowdown in productivity in recent decades.

- (3) In each country, it grows linearly up to 2100 to reach the highest level (across countries) reached at the end of the PWT sample.²² This implies no or little growth for the countries closest to that level in 2017, and a very quick catch-up for those further away (most of them). This scenario would be representative if we assume convergence, although possibly a stretch in many cases, including for several advanced economies.²³


Figure 1.12 plots the evolution of $W_{i,t}^{acc}$ for a selection of advanced economies under the three scenarios for $w^{H,nd}$, assuming a saving rate of 10%. Because of large differences in levels across countries and since we are interested in the relative increase or decrease in $W_{i,t}^{acc}$ over the years, this measure is normalised by dividing by the country's maximum observed $W_{i,t}^{acc}$ before the end of the sample, i.e. a value of 1 corresponds to the country's peak level of $W_{i,t}^{acc}$ up to 2017. To better see the trends, Figure 1.13 displays the simple average of each set of lines in Figure 1.12. A few facts appear evident from these figures:

- Thanks to robust technological and wage growth, wealth accumulation since 1950 has on average become steadily faster until around the year 1990, while it has declined significantly since 2000.
- Under the two more plausible scenarios for advanced economies, wealth accumulation in future decades is set to become harder or remain at current levels, which are lower than a generation ago but significantly higher than before the 1980s.
- Where $w^{H,nd}$ and $r^{k,nd}$ are both unchanged (scenario (1)), the impact of demographic dynamics on wealth accumulation through returns prevails on the impact through wages. In general, for most countries

²²The country with the highest $w^{L,nd}$ in 2017 was Norway.

²³For instance, the country in the 75th percentile of the distribution in the sample in 2017 had an $w^{L,nd}$ which was about 45% lower than Norway's, implying a growth rate well above (2) over the forecast horizon. The median country had an $w^{H,nd}$ 63% lower which implies an even higher growth rate.

Figure 1.12: *Wealth accumulated over 20 years from initial investment of 10% of wage, selected advanced economies.*



Note: Expressed as a fraction of its peak value up to 2017. The countries included are the same as for Figure 1.10.

across all three scenarios, in the absence of demographic effects $W_{i,t}^{acc}$ would either not decrease as steeply or it would increase more strongly (i.e., demographic factors weigh down on the growth of $W_{i,t}^{acc}$ in all three scenarios).²⁴


- Although unlikely, a return to growth levels last seen several decades ago in advanced economies could still lead to improving wealth accumulation prospects for future generations, despite demographic developments.

While the above findings relate to how one year's worth of saving can potentially accumulate, I now want to explore how different generations fare once their whole life-time wealth accumulation journeys are compared. Specifically, given changing wage and return levels, how does the wealth accumulated by a worker saving throughout their working life change across generations?

Consider a worker who starts working in year t in country i and initially

²⁴Results are available but omitted for brevity.

Figure 1.13: Wealth accumulated over 20 years from initial investment of 10% of wage, as a fraction of peak value, average across selected countries.



Note: Averages of each set of lines in Figure 1.12.

saves a proportion s_r of their after-tax income $(1 - \tau_w)w_{i,t}^L$, where τ_w is the tax rate on labour income. I indicate this amount as $S_{0,i,t}$, corresponding to total savings (wealth) at the end of the first year of saving when saving began at year t . This saving is invested and returns $s_r(1 - \tau_w)w_{i,t}^L \times (1 + (1 - \tau_k)r_{i,t}^k)$ at year $t + 1$, where τ_k is the tax rate on capital gains.²⁵ The worker adds additional savings equal to $s_r(1 - \tau_w)w_{i,t+1}^L$, such that total savings are now $S_{1,i,t} = (1 - \tau_w)w_{i,t}^L \times (1 + (1 - \tau_k)r_{i,t}^k) + (1 - \tau_w)w_{i,t+1}^L = S_{0,i,t}(1 + (1 - \tau_k)r_{i,t}^k) + (1 - \tau_w)w_{i,t+1}^L$. In general, for any $j \neq 0$:

$$S_{j,i,t} = S_{j-1,i,t}(1 + (1 - \tau_k)r_{i,t+j-1}^k) + s_r(1 - \tau_w)w_{i,t+j}^L \quad (1.17)$$


I set both τ_w and τ_k at 35%, and s_r at 20%²⁶, and I obtain an estimate of

²⁵Taxation rates were not introduced above but being constant they have no impact in terms of the direction of any movements in W^{acc} etc. The effect of demographics through changes in tax rates is likely to make effects shown in this Section even larger than displayed here, as a larger proportion of retirees tends to lead to higher public expenditure and therefore higher tax rates, lowering after-tax savings and rates of return. To keep the exercises simple, I do not add these dynamics but it should be kept in mind as a potential further effect.

²⁶Again, the tax rate numbers are indicative and based on representative values; changing them would have an impact on the values of W^{acc} , but not its dynamics over time. As to savings rates, it would be reasonable to assume, based on the extensive

$S_{39,i,t}$, corresponding to 40 years of savings, for each country and year based on each of the three scenarios for $w^{H,nd}$. To obtain Figure 1.14 I go through all the steps as for Figure 1.13, i.e.: (i) I compute $S_{39,i,t}$; (ii) for each country, I divide it by its peak pre-2017 value across all years; (iii) for every year, I take the average of these values among the countries in the sample used for Figure 1.12 (averaging all the countries in the sample returns similar results).

Figure 1.14: *Wealth accumulated over 40 years' working life, as a fraction of peak value, average across selected countries.*



Note: See text for computation method. The x-axis represents the year at which the agent starts saving (i.e. year t in equation 1.17).

A similar picture to Figure 1.13 emerges, where on average the peak in accumulated wealth is expected to be for the cohorts who started saving in the 1990s. Assuming they started working in their twenties, these agents are currently in their fifties and will likely retire starting from around year 2030. They are expected to have significantly higher accumulated wealth than the cohorts before them, but in advanced economies they will likely also be the last generation to experience a significant increase in wealth

literature on life-cycle consumption and savings, that savings rates are age-dependent, however here they are kept constant (with the implication that any arbitrary choice of s_t for the sake of this exercise will have no impact on the relative sizes of W^{acc} between cohorts). Nevertheless, I find that adding age-dependent savings rates to this exercise, based for instance on saving patterns like described in Feiveson and Sabelhaus (2019) [Feiveson and Sabelhaus \(2019\)](#), has only a negligible impact on results as shown below, leading to the same conclusions.

compared to their parents' generation. Depending on the forecast for $w^{H,nd}$, subsequent generations (i.e., people currently at work and younger than 50, and their future descendants) could experience (i) progressively lower wealth accumulation (most likely scenario for advanced economies with slow growth); (ii) broadly same level of wealth accumulation as their elders; or (iii) continuously rising wealth (most likely for economies furthest from the technological frontier). Based on this, material improvements in standards of living from generation to generation might not continue to be the case in advanced economies, while to an extent it is already so for some countries.²⁷

This is not all, however. A higher level of wealth at retirement does not guarantee higher standards of living when retired, for two main reasons:

- Lower returns on capital result in lower levels of income generated from any given level of accumulated wealth.
- Longer life expectancies mean that retirees need to make sure their wealth can sustain their livelihoods for a longer period of time, decreasing the pace at which they are able to decumulate it.

Because of these facts, even if current levels of wealth accumulation were to remain, in several countries future generations will likely not be able to maintain the same standard of living at retirement as their parents. In this last exercise, I estimate how much retirees can expect to consume out of their accumulated wealth by investing it and decumulating it, based on the rate of return on capital, their residual life expectancy, and their initial accumulated wealth.

Consider an agent aged 65 who retires at year t in country i . They started working and saving 40 years before, and have accumulated wealth equal to $S_{i,39,t-40}$. They have an expected residual life expectancy at 65 of x but they want to provision for a higher number of years in case they live longer, say $1.25x$. They want to generate a constant cash flow out of their wealth via a combination of capital returns (at rate $r_{i,t}^k$) and sale of assets, such that at year $t + 1.25x$ they have completely decumulated all of their wealth. This situation is equivalent to a hypothetical "reverse mortgage" where the agent loans out $S_{i,39,t-40}$ for $1.25x$ years at $r_{i,t}^k$; every year they receive a constant cash flow payment which includes principal repayments


²⁷Southern European countries being the most striking example in this case, based on these data. See below.

(i.e. the amount of asset sales) plus interest (i.e. returns on the remaining capital). Based on the typical formula for mortgage amortisation, the constant annual repayments they can expect to receive are equal to:

$$A_{i,t} = S_{i,39,t-40} \left(r_{i,t}^k + \frac{r_{i,t}^k}{(1 + r_{i,t}^k)^{1.25x} - 1} \right) \quad (1.18)$$

As previously, I estimate this for every country and year, and for each country I compute the annual value relative to the pre-2017 country peak, then average this for a subset of countries, as done above for Figure 1.14.²⁸ Figure 1.15 shows the results of this exercise, which are striking. See further Figure 1.16 for countries in Southern Europe, which are affected particularly hard by low productivity growth since 2000, and are seeing living standard prospects of retirees possibly reducing by a half by the end of the century.


Figure 1.15: *Annual cash flow from reverse mortgage of wealth accumulated over 40 years' working life, as a fraction of peak value, average across selected countries.*



Note: Country-year values for the cash flow were calculated as per equation 1.18; the lines displayed in the figure are calculated as for Figure 1.13 and Figure 1.14. The x-axis represents the year at which the agent retires.

²⁸Annual residual life expectancies at 65 for the purposes of this exercise are backward induced from the UN World Population Prospects projections. This method matches existing data (where available) with a sufficient level of accuracy, but may differ somewhat.

Figure 1.16: *Annual cash flow from reverse mortgage of wealth accumulated over 40 years' working life, as a fraction of peak value.*



Note: Country-year values for the cash flow were calculated as per equation 1.18. The x-axis represents the year at which the agent retires.

On average, agents currently retiring²⁹ with the wealth accumulated during their working lives can expect to enjoy some of the highest standards of living at retirement across all cohorts in the sample; for all generations currently at work and far from retirement, prospects are expected to get progressively worse. Even under the most optimistic scenario, standards for retirees are set to decline from current levels until around 2050, and under the first and second scenarios they are seen declining all the way to the end of the projection horizon, with different intensities.

1.7. CONCLUSIONS

The findings in this paper highlight stark inter-generational inequalities in ageing economies when it comes to wealth and its accumulation, with the cohorts born at the start of the demographic transition benefiting most from demographic dynamics and growing economies, while subsequent generations are set to lose. This may be part of the reason why younger generations in many advanced economies have been expressing increasing unease at their future prospects when comparing themselves with their

²⁹I.e., they would have been born in the late fifties and sixties.

parents' generations.

They may be trapped in this “inverted-pyramid scheme” which is the economy of an ageing society. Normal pyramid schemes can only be sustained as long as an exponentially increasing number of participants join the scheme; once this is no longer the case, the scheme fails. While, like in a pyramid scheme, economies in modern times had been relying on an increasing number of new “entrants”, which ensured rapidly increasing standards of living and wealth for all economic agents, positive demographic trends have now reversed in several economies, dragging down future prospects for generations to come. While in a pyramid scheme, those already in it benefit from new entrants (and every cohort but the last is set to gain), in a world of inverted population pyramids, new generations in the economy are increasingly penalised by the existence and size of the previous generations.

The implications of the results of this paper are clear: future generations in advanced economies can expect, under reasonable assumptions, to either accumulate less wealth (restricting their retirement income), or to save a significantly higher proportion of their labour income during their working age to avoid this (thus limiting consumption during their working age). In addition, depending on the social security system in place, they may either expect to pay higher taxes during their working life, or to receive lower pension income from the state when retired. In any case, it appears that they can reasonably expect to be worse off than previous generations as far as wealth accumulation is concerned.

While the changes in society and demographics that underpin these implications are in all likelihood here to stay, there are a number of policy actions that can help mitigate the adverse effects of the demographic dynamics we are facing.

- A return to strong technological growth, of the type experienced during the second half of the twentieth century (but more sustainable), would ensure that standards of living can remain constant or keep growing, even while populations age and decline. This would require large public and private investment in research and development and human capital which might however exacerbate the existing issues arising from abundance of capital.
- Increasing focus on productive vs unproductive assets. The increasing prices of residential capital, while other forms of capital have

been depreciating, probably suggest that there is over-investment in housing. This is also partly because of a lack of available opportunities for productive investments (see point above), but it goes both ways.

- Improvements in the levels of financial education of the population. Many households lack the financial education to confidently invest in riskier asset classes, and these households are usually also less wealthy to begin with. A better financial education of all households would allow a greater proportion of them to knowingly invest in riskier but higher-yielding asset types, which have typically been held by the wealthiest in society. It would also likely channel investment into more productive forms of capital.
- Increasing retirement ages, and especially allowing agents to work longer work lives as increased life expectancies mean they can also expect to stay healthy and productive for longer.

Several complex aspects and interactions between demographic developments and the macroeconomy, which might add further dynamics and insights into the effects through which population ageing affects wealth inequality, could not be covered within the scope of this paper, but warrant further study. These include (i) the impact of inter-generational transfers on the transmission of wealth inequality from generation to generation, as agents live longer and have fewer children; (ii) dynamics related to inequality *within* generations and how this translates into inter-generational inequality; (iii) portfolio allocations and returns across different asset types, and how this affects how agents with different portfolio allocations accumulate wealth; (iv) the impact of housing on reducing or increasing wealth inequalities. Many of these issues would be best analysed within the context of a general-equilibrium overlapping-generation model and are generally not suited for an empirical approach. This is left for future research.

2 Housing policies, homeownership and inequality

2.1. INTRODUCTION

For the majority of people, wealth is housing wealth. Housing is simultaneously one of the most important goods consumed by individuals, while also dominating most households' asset portfolios. For this reason, Government policy that alters the distribution of homeownership, or its price relative to renting, either by design or indirectly, will also carry implications for the distributions of household wealth and welfare. These policies can for instance operate by altering relative prices of housing and rental properties, and by altering the availability of financing across agents. While some policies are created with distributional objectives, hoping to improve welfare of one group relative to another, others can have other objectives, such as serving critical functions in financial and macroeconomic stability.

In this paper, we quantitatively estimate the impact of various policy measures in a heterogeneous agent life-cycle model with endogenous selection into homeownership. This setting allows us to estimate the consequences of a wide range of factors that affect the distribution of homeownership, income, wealth, and consumption. In particular, we study changes in: interest rates, supply side constraints, taxation on rental income, macroprudential credit constraints, and the presence of institutional investors. Some of these factors, such as the presence of institutional investors, are not policies per se, but are an important factor in the housing market, that can be affected by legislation. Others, such as macroprudential credit regulations, are policies whose benefit are difficult to capture in our experiment, but have large potential implications for the allocation of credit across the distribution. We therefore emphasise that the distributional consequences

of a particular policy should be taken as a partial story, and in the case of critical stability tools, may be a necessary reality that should be taken into consideration when sculpting related housing policy. We explore a number of combinations of policies, credit, and supply environments in tandem to understand whether negative distributional consequences of one might be offset by another.

Two papers that take a similar approach to understanding the role of housing policy are: [Alpanda and Zubairy \(2016\)](#) and [Alpanda and Zubairy \(2017\)](#). Both investigate the role that a set of government policies play in housing markets using general equilibrium models featuring three household types: savers, borrowers, and renters. [Alpanda and Zubairy \(2016\)](#) focuses on fiscal tax policies directed at housing and their output implications.¹ Using the same model, [Alpanda and Zubairy \(2017\)](#) instead focus on monetary and macroprudential policies and their implications for household indebtedness in conjunction with fiscal policies such as property taxes and mortgage interest deductibility. Importantly for our work, they find that there are differential effects across the three groups studied, discussing the potential redistributive natures of these policies. Our model expands on this heterogeneity by making the entrance of agents into the housing market endogenous, allowing for a more credible study of the nature of such redistribution.

The distributional impacts of macroeconomic policy more broadly are a growing concern. A large literature has grown from work such as [Doepke, Schneider, and Selezneva \(Doepke et al.\)](#) that explores the role that traditional monetary policy instruments have on distributions of wealth. Traditional policy operates through wealth channels as inflation alters the real value of assets, redistributing wealth among net borrowers and lenders in the economy. Their work and others² also investigate the role of *unconventional* monetary policy tools such as quantitative easing. Quantitative easing alters interest rates and credit availability, similarly redistributing wealth across these groups. Relevant to our question, [La Cava and He \(2021\)](#) show that conventional interest rate policy may have distributional consequences that operate through local housing markets due to the greater sensitivity of high price localities to these interest rate shocks. Recent work by [Amberg et al. \(2022\)](#) identifies five potential channels for traditional

¹Specifically they find large negative output implications for raising tax revenue through housing policy, with a long-run multiplier of around 2.

²See, for example: [Hohberger et al. \(2020\)](#), and [Davtyan \(2017\)](#)

policy highlighting a number of potential channels, and provides a recent and useful overview of these potential mechanisms.

In addition to conventional housing policies we also study *macroprudential* policy. These have distributional effects because not only is housing a larger component of wealth among lower and middle class households than at the top of the distribution,³ but common macroprudential tools in the form of loan-to-income (LTI) and loan-to-value (LTV) ratios explicitly create barriers tied to existing income and wealth. Moreover, rental properties are an important source of income for those at the top of the distribution while serving as a large expense for the bottom, and any impact of credit availability on relative prices of bought and rented housing can alter both the income and consumption inequality across these groups. For this reason, we study the role of a range of policies and macroeconomic conditions across loose and tight credit regimes, and how these affect inequality, acknowledging that such regimes may serve important roles in improving welfare through financial stabilisation.

The aftermath of the Global Financial Crisis saw an explosion of macroeconomic research on the role of housing in the macroeconomy. Largely, this literature focused on cyclical debt dynamics and the outsized role that housing markets can play in financial booms and busts. These papers use heterogeneous agent models and study the role that credit constraints play on macroeconomic stability and the housing market. In doing so, [Kaplan et al. \(2020\)](#) find weak evidence for the transmission of credit constraints to house prices, while [Greenwald and Guren \(2021\)](#) suggest a more important link between the two. Both the level of housing costs and the relative cost of buying a house vs renting matter for the distributional impacts that we study. Our choice of modelling housing supply through a stylised construction sector will pin down the price of housing in a way similar to [Kaplan et al. \(2020\)](#), but allowing for endogenous choice to fulfil housing needs via ownership or renting will allow policies to affect the *relative* cost of renting and housing. We build from this literature, using similar modelling tools to address distributional outcomes. In [Greenwald and Guren \(2021\)](#), the key determinant of the price-rent ratio's response to credit is the response of supply both through the construction sector, as well as an important "tenure supply" channel, in which landlords supply owner versus rented housing stock. This means that in the presence of

³For an empirical exploration of this and other stylised facts across OECD countries, see [Causa et al. \(2019\)](#).

highly segmented rental and ownership markets, credit conditions may have quite large impacts on house prices.

In quite complementary recent work to ours, [Castellanos et al. \(2024\)](#) show that landlord entry and exit can additionally create a channel for credit conditions to alter home prices and therefore inequality as rising rent/price ratios push more poor young households into rental markets while benefiting older and richer households. Like [Castellanos et al. \(2024\)](#), we show that credit constraints tend to decrease homeownership rates, and like them we show that this also tends to increase inequality in both consumption and income, operating largely through an increase in the rental yield relative to house prices. Studying these policies while increasing the supply of housing⁴ actually can further reduce homeownership rates, with the excess supply bought by high income landlords, but this in turn substantially lowers inequality of consumption due to much lower rental costs for the (poorer) renters. In addition to the authors' findings on the distributional effects of borrower-based macroprudential policies, which we broadly confirm (while also separating deposits and loans as two different asset classes in households' portfolios), we also analyse a number of other housing market and related policies, and their distributional effects, expanding the scope of the analysis and findings.

Indeed, we find many combinations of policies and credit conditions that may increase inequality in terms of homeownership while reducing inequality in consumption, and by association the welfare of the bottom of the distribution relative to the top. While a host of policies aimed at improving homeownership exist, there is little work that quantifies their distributional impact. The value of heterogeneous agent modelling is that it allows us to understand the distributional role that policies play through transferring income from one group to another largely through changes in relative prices of these different classes of housing. In an investigation of Britain's help-to-buy scheme, [Carozzi et al. \(2024\)](#) show that such policies largely work to increase prices and improve the financial performance of developers without helping buyers enter the housing market.

[Dustmann et al. \(2022\)](#) show that income inequality is amplified when looking at income post housing expenditure. Their work documents this trend in Germany over the last three decades, showing that relative

⁴As we explain in detail below, we consider either productivity of the construction sector or the availability of land permits as a proxy for government supply-side housing policies.

cost of ownership to rentals have driven a large share of this difference. Their work shows for housing a phenomenon, documented by [Kaplan and Schulhofer-Wohl \(2017\)](#), of lower income households experiencing higher inflation in the basket of goods they consume. [Christophers \(2021\)](#) discusses the need to understand how inequality in housing tenure, i.e. between renters and buyers, explains trends in housing wealth inequality. Work such as [Bartels and Schröder \(2020\)](#) show that in Germany from 2002-2017 a large contribution to increases in income inequality has been that of rental income. This is consistent with the role that rental markets play in [Kindermann and Kohls \(2018\)](#), who use the ECB's *Household Finance and Consumption Survey* to study the role of *wedges* in rental markets⁵ in determining wealth inequality. A broad read of the literature on inequality across housing, income, wealth, and consumption ultimately points to consistent, but complex relationships between these. For this reason, we are interested in all four of these forms of inequality. From a welfare perspective, consumption might be the most salient, but undoubtedly policymakers are interested in housing inequality in itself, together with its relationship to income and wealth distributions. It is important to know how trade-offs between different types of inequality may arise; we see a few examples of this play out in our results.

[Causa et al. \(2019\)](#) show that there is a strong negative correlation between homeownership and wealth inequality in OECD countries. Removing housing from wealth increases Gini coefficients among these countries on average by about one-quarter, suggesting that any policies that reduce rates of homeownership could be a channel for increased wealth inequality. Existing empirical work shows that macroprudential rules have the potential to have distributional consequences. [Georgescu and Martín \(2021\)](#) use survey data on housing and consumption within the euro area to show that LTV credit limits and debt-service-to-income limits moderately increase wealth inequality. Their work attempts to simulate both the negative impact caused by constraining households, as well as the positive effects of reducing the probability of financial crises. In an agent-based model, [Tarne et al. \(2022\)](#) show how it is critical to differentiate between which agents are affected by these credit restrictions. Restrictions on *buy-to-let* investors can have positive effects on equality while such

⁵In their life-cycle model this wedge is akin to an iceberg trade cost where only a fraction of the housing units that landlords supply arrive at the potential tenant; this is a stand in for what they call “institutional frameworks” that affect rental markets, but could also stand in for policies that make renting the same housing unit relatively more expensive.

restrictions on first-time buyers works in the opposite direction. [Acharya et al. \(2020\)](#) study the effects of LTI/LTV limits and find that mortgage credit is reallocated from low income to high income borrowers, with spacial effects from urban to rural counties. This relieves price pressure in “hot” markets. The banks reallocate their portfolio to take more risks in other securities and corporate credit. In the Irish mortgage market, which serves as a useful environment to study these rules, [Kingham et al. \(2019\)](#) find that the introduction of LTI/LTV rules induced wealthy borrowers to increase their down-payments, while poorer households tended to shift toward lower value housing stock. At present our work finds little impact on wealth inequality itself, which moves little in part due to financial market participation being counterfactually high,⁶ however we do find substantial decreases in homeownership and increases in other forms of inequality.

A small, but growing, literature has studied the impact of macroprudential regulation on inequality in the context of quantitative life-cycle models, seeking to estimate the welfare cost of these policies. [Mavropoulos and Xiong \(2019\)](#) use panel data and a structural life-cycle housing choice model to show that LTV/LTI worked to limit home-ownership and wealth accumulation in the post crisis era in Europe. Interestingly, while they find large decreases in wealth accumulation due to the policy, they are welfare improving due to decreases in risk. [Rubio and Carrasco-Gallego \(2014\)](#) do a welfare analysis of LTV ratios in a DSGE model. Their model includes two agent *types* representing borrowers and savers, with LTV tightening harming borrowers while benefiting savers. Work such as [Mavropoulos and Xiong \(2019\)](#) finds that macroprudential policy improves welfare on net, driven largely by the well documented positive effects of macroprudential effects on housing risk, which enter into their partial equilibrium framework exogenously and affect preferences of households directly through the reduction in foreclosure risk. While this seems reasonable, it simplifies the boom-bust cycle that is documented in [Kaplan et al. \(2020\)](#) and others.

Our work contributes to this literature by quantifying the effects of LTI and

⁶Extensions of our model that limit the ability of households to accumulate non-housing wealth by including, for example, costs of accessing financial markets might help connect our results more closely to the literature and data on the connection between housing wealth and inequality. In the present form we have avoided adding an additional layer of *ad hoc* heterogeneity to keep from muddying our understanding of the model results.

LTV policies on the distributions of (housing and total) wealth, income and consumption. Because we are able to provide insights into the mechanisms that drive this effect, we are also able to provide insight into the potential policy options available to offset these distortions. We find that tightening loan-to-income ratios to 3.5 (from 6) substantially reduces homeownership by the bottom 50% of the wealth distribution, having a smaller but still sizeable impact on their share of overall wealth. Gini coefficients for wealth, housing, income, and consumption all increase with such credit constraints.

We also explore the distributional consequences of a few other policies that have received less attention in the literature, and find compelling results. We highlight especially how supply can greatly alleviate or aggravate inequality caused by other policies or features of the housing market. Our treatment of the construction sector is relatively simple, but we are still able to capture the overall effect of policies that support the expansion of housing supply, such as increases in land permits or an improvement in the planning process. These lead to lower house prices overall, and as we will show below, these lower prices also come with lower rental yields and thus much lower post-housing income inequality. Perhaps counter-intuitively, a decrease in house prices does not always lead to increased homeownership, as there are rent levels low enough at which some households can decide that they are better off renting; nevertheless, welfare and consumption are always significantly more equal when house prices are lower, and for this housing supply is fundamental.

Tax policies can also have large effects on how income is redistributed between renters and landlords. In the case of rental income tax, we note how, by creating a wedge between the rent paid by renters, and the net income received by landlords, such a tax is harmful for both categories (with only owner-occupiers unaffected by it). Within our model, we quantify that the burden of the tax actually falls mostly on renters, whose rental payments rise due to the tax much more than landlords' net returns fall. Renters being also much poorer than landlords to begin with, they are hit even more relative to their income. These dynamics would thus paradoxically make a tax levied on landlords highly unequal and unfair on tenants.

The inclusion of foreign institutional investors to our baseline model greatly changes the dynamics of the housing market. By essentially competing against each other as well as the domestic landlord households for rental

investment properties, they can lead to significantly lower rental yields. This is great news for tenants, who enjoy higher disposable income and consumption, but bad for domestic landlords. Needless to say, inequality in the household sector decreases as a result, but at the expense of lower homeownership rates, as well as lower overall wealth (since a significant chunk of it ends up abroad). While in a number of cases the positive effects from the fall in inequality far outweigh the effect of declines in homeownership, leading to general improvements in welfare, these outcomes are however dependent on adequate regulation, ensuring these entities act competitively and do not end up 'cornering' the housing market. On the other hand, we note how rental housing provided by the Government is conceptually similar in nature to institutional investors, and while it can provide some of the same benefits, it can potentially also avoid some of the costs. Coupled with tight macroprudential policies, the presence of institutional investors can lead to even lower homeownership rates in the population.

The rest of the paper is structured as follows: Section 2.2 describes our modelling framework; Section 2.3 outlines the calibration of the relevant parameters of our model, and the solution algorithm; Section 3.4 discusses the results of the various simulations of the model, and Section 2.5 concludes.

2.2. MODEL

We construct a heterogeneous agent life-cycle model of housing choice, with a housing and rental market. Using this model, we study the distributional implications of the change or introduction of a number of policies or characteristics of the market, with the aim of evaluating their impact from a point of view of inequality. In this section, we describe the structure and main assumptions of this model.

2.2.1. Households

Demographics and Preferences

Finitely lived households survive for a maximum of 80, annual, periods, with conditional survival probability at each age, j , given by ς_j . They work for wages during the first 45 periods of life, at which point they retire. Households maximise their lifetime utility over consumption of goods and

housing services given by:

$$\mathbb{E}_0 \left[\sum_{j=1}^N \beta^{j-1} \varsigma^{j-1} u(c_j, s_j, a_{j+1}) \right] \quad (2.1)$$

where β is the discount factor, c_j household consumption of non-housing goods, s_j is their consumption of housing services, and a_{j+1} is their net wealth at $t + 1$. Preferences are such that the instantaneous utility function u is given by:

$$u(c, s, a) = \frac{[(1 - \phi)c^{1-\gamma} + \phi s^{1-\gamma}]^{(1-\theta)/(1-\gamma)} - 1}{1 - \theta} + (1 - \varsigma_j)\psi(s_{t+1}, l_h) \quad (2.2)$$

where the first part represents the utility from consumption and housing services; the elasticity of substitution between consumption and housing services is given by $\frac{1}{\gamma}$ and the intertemporal elasticity of substitution is $\frac{1}{\theta}$. The parameter ϕ measures the taste for housing services relative to consumption goods. Households are divided into two types, which differ solely in their preference for housing. Transfers of bequests happen only between households of the same type.

The additive ψ factor represents a “warm-glow” bequest motive similar to that used in [De Nardi \(2004\)](#), however unlike [De Nardi \(2004\)](#), households derive a warm glow from bequests of housing wealth (and only from the first housing unit if any). Specifically:

$$\psi(s_{t+1}, l_h) = \psi_1 \left(1 + \frac{l_h}{\psi_2} \right)^{1-\theta} \quad (2.3)$$

where l_h is an indicator equal to 1 if $h > 0$. This encourages households to remain homeowners, which is useful in our case as households in the model otherwise liquidate their lumpy housing assets to smooth consumption in retirement resulting in counterfactually low levels of homeownership in older age groups.

Housing Services

Consumption of housing services comes in one of two broad types, ownership and renting. Denote r_h the rental rate for housing services, and p_h the price of a unit of housing purchased. Renters receive $s(h = 0) = 1$ in home

services while homeowners $s(h > 0) = \omega > 1$, providing a linear premium in home ownership associated with the first house. This provides a utility incentive to own rather than rent, but only for the first house owned (i.e., owning more than one house does not provide any additional utility to the owner, only potential financial gain from rental income). We assume that homeowners occupy their first unit, $h = 1$, of housing and rent remaining units, $h > 1$ to other households.

Homeowners can purchase housing outright or through a mortgage. Mortgage holders borrow money from banks to finance the purchase of their homes and make interest payments each period. Renters pay rental payments to landlords corresponding to a fraction r_h (the rental yield) of the house price p_h . All homeowners, with mortgage or not, face costs associated with the depreciation of their home equal to $\delta_h p_h$, as well as property taxes $\tau_h p_h$ for each housing unit they own.

Financial Markets

Households have access to a risk free bond, b , which pays r_b one period in the future. These bonds are offered by externally-owned financial institutions, with r_b determined exogenously. Bond interest income is taxed at rate τ_k . This is the savings vehicle of households, who cannot borrow in bonds. However, they can access mortgage borrowing (also from the same external, deep-pocketed financial institutions) to finance the purchase of housing. The interest rate on mortgage loans is equal to $r_l = r_b + \kappa$, where κ is a exogenous premium earned by the financial intermediaries who provide loans and safe assets in the form of bonds.

Borrowing is limited by macroprudential lending regulation and thus subject to maximum LTI and LTV limits. Therefore a household's maximum level of borrowing is subject to the following two constraints:

$$\begin{aligned} d_j &\leq \lambda_{LTV} p_h h_j \\ d_j &\leq \lambda_{LTI} y_j \end{aligned} \tag{2.4}$$

Where λ_{LTV} and λ_{LTI} denote the macroprudential loan-to-value and loan-to-income ratios respectively. Further, when households apply for a mortgage they cannot exceed retirement age, and the loan needs to be below a certain fraction of their expected remaining lifetime labour income.

At each new period, households with debt need to pay $r_l\%$ of their loan as interest repayments. The loan principal can then be repaid in full, in part, or can be refinanced (subject to remaining within the macroprudential and other limits).

Income process

Households' labor income is given by $W_t \ell_{j,t}^i$. Where W_t is the aggregate wage and households are subject to idiosyncratic productivity shocks given by:

$$\ell_{j,t}^i = \exp(\varepsilon_{j,t}^i) N_{j,t}^i \quad (2.5)$$

The term $\varepsilon_{j,t}^i$ is a transitory, mean zero, shock realised in each period, while the permanent component of households' wages follows an age specific trend $\eta_{j,t}$ given by:

$$N_{j,t}^i = N_{j-1,t-1}^i \exp(\eta_j) \quad (2.6)$$

2.2.2. Government

The Government taxes labor income, interest income, and property wealth, with the respective marginal tax rates being τ^ℓ , τ^k and τ^h . It may also levy a tax τ^{rh} on rental income (this is one of the policies we look at later), i.e. each rented house will generate revenue corresponding to a fraction τ^{rh} of the rental income $r_h p_h$ it generates. Government revenues are thus given by:

$$G_{Rev} = \tau^\ell W_t N^w + \tau^k B_t r^b + \tau^h p_h N_t (+ \tau^{rh} r_h p_h R_D) \quad (2.7)$$

where W_t is total wages earned in the economy, B_t is total financial assets (excluding debt), and R_D is the total demand for rental properties i.e. the number of properties earning rental income. These revenues are partly used to fund pension income ($Tr(\ell_{J_R,t}^i)$) for retirees; the individual pension income of a household is heterogeneous and depends on labor income in the final year of working life. Government transfers to individual (retired) households are thus given by:

$$Tr(\ell_{J_R,t}^i) = \gamma_{repl} \ell_{J_R,t}^i \quad (2.8)$$

which says that the Government provides pension income to retirees, in proportion to their final working income; γ_{repl} represents the replacement rate of pension payments relative to an agent's income during their last year of working life. Typically, the Government will spend less than its total revenues on pensions; we assume that the rest is spent on public services providing a certain level $u(g)$ of utility to all agents equally. This means that the Government runs a balanced budget, and also that $u(g)$ can be ignored as it does not influence households' utility maximisation problem.

2.2.3. Production

The production of final consumption goods follows a constant returns to scale production function, with labour, the only input, equal to the number of households of working age:

$$Y_t = A_t L_t \quad (2.9)$$

where Y_t is total output, A_t is the aggregate labour productivity, and L_t is the labour supply, which is also equal to the number of agents in working age, N_w . As a result, the aggregate wage level is equal to $W^L = Y_t/L_t$ (pre-tax), with individual wages differing as per above.

Construction Sector

Housing units are built by a construction sector, which produces Y_h units every year. The simulation starts with one unit of housing per capita, and to maintain this the construction sector needs to meet demand for new housing from the new population (equal to the number of newly born households minus the deceased: $n_t N_{1,t-1} - \sum_{j=1}^{J-1} \varsigma_j N_{j,t-1}$), plus maintenance of existing properties, which depreciate at rate δ_h . We assume that depreciation forces homeowners to pay for repairs costing $\delta_h p_h$ for each unit they own, leading to a demand for construction equal to $\delta_h H_s$, where H_s is the total housing supply.

Given depreciation and construction output, the housing supply changes according to the following law of motion:

$$H_{s,t} = (1 - \delta_h) H_{s,t-1} + Y_{h,t} \quad (2.10)$$

To keep up with population growth and depreciation (there needs to be at least one housing units for each household at all times), for an equilibrium we have the requirement that:

$$Y_h \geq n_t N_{1,t-1} - \sum_{j=1}^{J-1} (1 - \varsigma_j) N_{j,t-1} + \delta_h H_s \quad (2.11)$$

The production function of the competitive construction sector is Cobb-Douglas, with inputs land L , and materials M :

$$Y_h = A_h L^\alpha M^{(1-\alpha)} \quad (2.12)$$

where A_h is the productivity of the construction sector, and α is the constant share of land in production. The amount of land permits is fixed to $\bar{L}$ and determined by the Government. The price of a unit of land is p_L (determined competitively), while we set the price of materials to 1. Profit maximisation (with respect to the amount of materials) among competitive construction firms results in the following housing investment function:

$$Y_h = A_h^{(1-\alpha)} ((1-\alpha)p_h)^{(1-\alpha)/\alpha} \bar{L} \quad (2.13)$$

with construction increasing with prices. An equilibrium between demand and supply of housing will only be reached if the supply of housing is exactly equal to the number of households N_t , i.e. at any period the construction sector should produce exactly as many units as necessary to keep up with depreciation and population growth, and no more. This is achieved if:

$$Y_h(p_h, \dots) = n_t N_{1,t-1} - \sum_{j=1}^{J-1} (1 - \varsigma_j) N_{j,t-1} + \delta_h H_s \quad (2.14)$$

where there is only one price, p_h^* , that allows for this. In other words, in a steady state, the house price is determined by the construction sector.

2.2.4. Institutional Investors

In a number of simulations, housing units can also be purchased by institutional investors, who are deep-pocketed entities with funding from

abroad whose business model is to purchase housing to generate rental returns. Whether or not they decide to invest in the domestic housing market depends on the level of rental returns that they can achieve; in the absence of aggregate risk, the house price does not matter to them. We assume that these investors face a certain cost of investing into the local housing market. This is given by some external funding costs that they face.

If the net domestic rental return (i.e., r_h net of rental income taxes, or $(1 - \tau_{rh})r_h$)⁷ is greater or equal than their funding cost (denoted by r^{ii}), plus any additional costs of owning a housing unit, i.e. depreciation and property tax, then the institutional investors will demand as many housing units as are available for them to purchase, as local housing becomes profitable for them. r^{ii} depends on external factors and is set exogenously.

We assume that institutional investors are limited in the number of housing units available for them to purchase. Specifically, households have first access to housing units, i.e. institutional investors can only purchase any units corresponding to the difference between the housing supply H_s and the household demand for housing H_d^{hh} .⁸ Therefore, the demand for housing by institutional investors, H_d^{ii} , will be:

$$H_d^{ii} = \begin{cases} 0, & \text{if } (1 - \tau_{rh})r_h < r^{ii} + \delta_h + \tau_h \\ H_s - H_d^{hh}, & \text{if } (1 - \tau_{rh})r_h \geq r^{ii} + \delta_h + \tau_h \text{ and } H_s > H_d^{hh} \\ 0, & \text{if } (1 - \tau_{rh})r_h \geq r^{ii} + \delta_h + \tau_h \text{ and } H_s \leq H_d^{hh} \end{cases} \quad (2.15)$$

By demanding any excess units not demanded by households for each level of r_h greater or equal to $(r^{ii} + \delta_h + \tau_h)/(1 - \tau_{rh})$ (call this r_h^{ii}) - and thus ensuring that the market for housing always clears in these cases - institutional investors effectively render r_h^{ii} the upper bound level of rental yields. In equilibrium, with institutional investors present, we have thus

⁷Note that if these investors are set up as corporations, they would pay tax on profits rather than on the rental income itself, i.e. tax would be paid after accounting for costs, and at a lower marginal rate. We abstract from this complication here but note that we also did not deduct mortgage interest payments from the rental income subject to tax in the case of households (as would be the case when calculating rental income tax) to avoid unnecessary complications. The fact that these investors may have a more favourable tax situation due to the lower tax rate on profits can instead be approximated by assuming a lower r^{ii} .

⁸If this were not the case, they would either own all or none of the housing units.

that:

$$r_h = r_h^{ii} = (r^{ii} + \delta_h + \tau_h) / (1 - \tau_{rh}) \quad (2.16)$$

2.3. SOLUTION AND CALIBRATION

We solve for an equilibrium that consists in a set of prices (w, p_h, r_h) such that after solving their maximization problem, households' demand for housing units is equal to the supply of housing, and the demand for rental units is equal to the rental supply. At the steady-state equilibrium, prices remain constant between periods; the population, output, housing supply and wealth grow at a constant pace.

To solve the model, we simulate it on a large number of households, and run it for a sufficient number of periods to allow for convergence to a steady state.

The procedure to determine an equilibrium is as follows:

- (1) We start from an initial guess of p_h and r_h
- (2) We estimate the optimal choices of households at every possible combination of age, wealth, housing, income, and debt, at these prices. The household optimisation problem is solved by value function iteration.
- (3) Based on the starting values of p_h and r_h , and households' optimisation, we simulate the model for a large enough number of periods until a steady-state is reached.
- (4) If at the steady-state, the housing and rental markets are not in equilibrium, we update p_h and r_h accordingly and repeat the process from (1); else, we have reached the steady-state equilibrium.

The steady-state solution of the model rests in the interactions between the market for housing units, the rental market, and the construction sector. For a solution, we have that

- The demand for housing needs to be at least equal to the population ($H_d \geq N$) to allow for every household to live in a unit, either rented or owned. We do not allow for the existence of homeless households.
- For the same reason, the supply of housing also needs to be at least equal to the population ($H_s \geq N$).

- The number of extra housing units owned by landlords (and/or institutional investors if they are a feature) must be at least equal to the number of households who do not own any housing units ($R_s \geq R_d$)

Additionally, as detailed above, there is no reason for households to demand more housing units than the population, as that would create empty housing with a cost to the owner but no benefit in terms of housing services to any tenant or rental return to the landlord. Accordingly in equilibrium:

$$H_s = H_d (= H_d^{hh} + H_d^{ii}) = N$$

$$R_s = R_d$$

where the rental supply is equal to all the housing owned by households in excess of the first unit (which is occupied by the owners N_{oo}) (plus all units owned by institutional investors, if present); rental demand is equal to the number of households who are not owner-occupiers.

$$R_s = H_d^{hh} - N_{oo} (+ H_d^{ii})$$

$$R_d = N - N_{oo}$$

This set of equations taken together tells us that whenever the housing market is in equilibrium (with housing supply also equal to the population), the rental market will clear as well.

2.3.1. Calibration

For the calibration of the parameters used in the model, we draw from the literature, or alternatively choose parameter values so as to match real-life data for Ireland; these data are derived primarily from the 2021 wave of the Household Finance and Consumption Survey (HFCS) ([Household Finance and Consumption Network, 2022](#)). Table 2.1 gives an overview of all the parameters used in the model, their values in the baseline model, and their source. Here we present briefly the rationale for the internally calibrated parameters.

The discount rate β , households' preferences for homeownership ω_1 and ω_2 , and the luxury of housing bequests ψ_2 were calibrated jointly to achieve a homeownership rate of 70% in the baseline scenario, with consistently high homeownership for retirees, as per the data; the resulting values

Table 2.1: Calibrated parameters of the baseline model.

Parameter	Name	Value	Source/Calibration
<i>Household utility</i>			
Discount factor	β	0.96	Internal
Preference for home-ownership (hh type 1)	ω_1	1.60	Internal
Preference for home-ownership (hh type 2)	ω_2	1.30	Internal
Share of households of type 1		0.6	
Preference for housing w.r.t. consumption	ϕ	0.30	
Elast. of substitution, housing and cons.	$1/\gamma$	0.8	Kaplan et al. (2020)
Inter-temporal elasticity of substitution	$1/\theta$	0.5	Kaplan et al. (2020)
Altruistic motive for bequests	ψ_1	-9.5	De Nardi (2004)
Luxury of beq. of 1st home	ψ_2	10	HFCS, Internal
Average prob. of being bequest recipient	π_{beq}	0.05	
Maximum number of housing units	h_{max}	4	HFCS, Internal
<i>Banking sector and mortgage market</i>			
Exogenous risk-free rate on deposits	r_b	0.00	Ireland
Intermediation wedge	κ	0.02	Ireland
Loan-to-income macroprudential threshold	λ_{LTI}	6	Policy variable
Loan-to-value macroprudential threshold	λ_{LTV}	0.9	Policy variable
<i>Construction sector</i>			
Land permits per capita	L	$(nN_t)/10$	HFCS, Internal
Productivity of the construction sector	A_h	1.3	HFCS, Internal
Depreciation of housing stock	δ_h	0.01	Standard
<i>Wages</i>			
Age-earnings profile: coeff. on age	β_{age}	0.0326379	HFCS, regression
Age-earnings profile: coeff. on age squared	β_{agesq}	-0.0004863	HFCS, regression
Age-earnings profile: constant	β_{const}	9.614493	HFCS, regression
Standard error of the white noise process	σ_ϵ	0.12	
Autoregressive coefficient	λ	0.985	
<i>Government</i>			
Tax on labour income	τ_w	0.34	HFCS, Ireland
Tax on interest income	τ_k	0.34	Ireland
Tax on property wealth	τ_h	0.001	Ireland
Tax on rental income	τ_{rh}	0.40	Ireland
Social security replacement rate	γ_{repl}	0.50	

are 0.96 for β , 1.3 and 1.6 for ω_1 and ω_2 , and 10 for ψ_2 . The HFCS also suggests that only very rarely do households own more than 4 housing units, so we set h_{max} at 4.⁹ We have two household types, differing by their levels of preference for housing (ω_1 and ω_2), because in real-life data the distribution of wealth appears to be bimodal, with the largest peak around 0 net wealth; in order to achieve this in the model, we need two types of households, of which one has less strong preferences for housing and is therefore less inclined to save:¹⁰ we set the share of households with this characteristics (type 1) at 60% of the population, giving a higher peak of the bimodal distribution at a low net wealth as per the data.

The age-earning profile of labour income was determined using the following regression on HFCS data:

$$y_i = \beta_{const} + \beta_{age} * a_i + \beta_{agesq} * a_i^2 + \gamma X_i + \epsilon_i \quad (2.17)$$

where y_i is the household's log labour income, a_i is the age of the reporting person minus 20 (to match our model), X_i is a vector of controls including household size, type, number of dependents and education levels, and we filtered for households whose reporting person is male, aged between 20 and 65 (i.e. with a_i between 0 to 45), and employed. The resulting age-earning curve was then scaled down such that, conditional on the steady-state population distribution, the average wage results equal to 1.

Still from the HFCS, we have that households with the lowest education level earn about three times less than those with the highest. We use this fact to determine the distance between the highest and lowest wage levels at each age, and we select 5 total skill levels. We made the transition between skill levels unlikely; this resulted in a standard deviation of 0.12 and an autoregressive coefficient of 0.985 for the idiosyncratic wage process. The creation of the respective Markov chain of transition probabilities between wage/skill levels follows [Tauchen \(1986\)](#).

The wedge between the interest rate on deposits and that on mortgages, κ , is set at 2%: according to Central Bank of Ireland data and our calculations, the average difference between mortgage and deposit rates since 2003 is just above 2%, while it has over time moved as low as less than 1% and as

⁹The number of housing units owner by any single household does not ever go above this number in our simulations even when left unrestricted.

¹⁰We could also have achieved this by having different values for β , for instance.

high as 3%. The same data tells us that nominal deposit rates have been well below 2% on average, with an ECB target inflation rate of 2%, i.e. a negative average real return; as our model is entirely in real terms, we select r_b to be 0%. We further select LTV and LTI thresholds of 0.9 and 6, respectively, as internal thresholds a prudent bank would set for itself in the absence of an external authority setting macroprudential policy.

For the construction sector parameters, A_h and L , we calibrate them jointly so as to achieve a steady-state equilibrium price of housing p_h that is consistent with the real-life average house price-to-income ratio in Ireland. Depending on how they are calculated (e.g. using only labour income vs all income, or only employees, or taking the median vs the average etc.), price to (gross) income ratios for Ireland, based on the HFCS, generally range from 5 to 7, and we settle for 5.8 in our baseline. We allow land permits L to expand with the population, and set them at $L = (nN_t)/10$, i.e. one every 10 new agents in the population;¹¹ the corresponding value of A_h giving us an equilibrium p_h of 5.8 is 1.3.

In terms of the tax rates, these were determined as follows: τ_k was set equal to the Deposit Income Retention Tax (DIRT), which stands at 34%; τ_w was determined as the average income tax rate theoretically paid by households in the HFCS sample when including all labour-related taxation plus credits, and this also stands at 34% based on our calculations;¹² τ_h is the approximate local property tax rate on a house worth €300,000, based on current LPT bands, i.e. 0.1%; and the property income tax is based on the higher marginal income tax rate of 40%.

On top of calibrating parameters of the model on Irish data, we also check that the results of the model match some distributional properties of the Irish data. We observe that, like in the data, the distribution of household wealth has a peak at zero, another (lower) peak around the median house price, and a long right tail. Due to there being only one house price in the model (as opposed to a different price for every property in real life), the second peak in the wealth distribution coming out of the model is much more pronounced around the common house price, but the rest of the distribution satisfactorily represents the original data. The model also does a decent job at reproducing the age distribution of tenancy, and

¹¹For the population itself, we take survival probabilities from the [Human Mortality Database](#) ([Human Mortality Database](#)) (Ireland, 2017), and set a cohort growth rate n of 0.5%.

¹²Note that while the system is progressive in Ireland, we keep one flat tax rate in our model to simplify things.

rates of homeownership with and without mortgage, as observed in the 2022 Irish Census data; just like in the data, most young households start as tenants and then homeownership rates increase as they age, first via mortgages (peaking in middle age) and then progressively owning housing outright (as households pay down mortgages), with retired households almost exclusively being homeowners.

Finally, we address certain choices made when setting the assumptions and parameters in the model. Firstly, households in the model have only one size (i.e. 1), which does not change over time (which is typical in overlapping-generations frameworks), while in real life the size of households changes across both cohort and age. The average size of households in the HFCS (in terms of consumption units) rises from about 1.5 (people) in the 20s, to 2 by age 50, falling gradually to 1.2 by 85; overall, the standard deviation of consumption units is about 0.6 with mean 1.7. While in the model the size of households remains constant (a household there can be conceptualised equally as a couple or as an individual), it can be shown that allowing households in the model to merge and split does not significantly change most of the distributional considerations we are looking at, while complicating the computation considerably. Household size matters however should we want to specifically compare, for instance, the situation of single individuals versus couples when faced with the same borrowing constraints and house prices, with couples benefiting from being able to pool resources. We leave these questions for future research.

Secondly, there is large disparity in house sizes/quality in the real world, which is in stark contrast with our assumption of one only size of housing with the same identical price. Among households that own their main residence but no other real estate, the average housing wealth is close to €300,000, with a very high standard deviation of almost €190,000, meaning that households do care about the size/quality of their house, and we are not able to capture this in our model. About 10% of the main residences in the data is worth over twice the median house, which is itself worth more than twice as much as another 10% of houses. While some models in the literature have incorporated house size considerations (see for instance [Castellanos et al. \(2024\)](#)), these however still do not nearly match the variation observed in the data due to significant computational difficulties; it is also due to these difficulties, and considering that any intermediate solution would not bring us any closer to the situation in the real world,

that we choose to abstract from this dimension completely.

Finally, the choice of not including risky assets, such as equity, in the model, while dictated by the complexity of the simulation and priority given to housing, is also largely borne out by the data. Participation in risky financial assets in Ireland is low, and the level of wealth held in these asset classes is only a fraction of real estate wealth. Only 10.5% of the households in the latest HFCS wave owned any publicly traded shares, while 3.3% owned mutual funds. The average holding of risky financial assets (publicly traded shares or mutual funds) was just below €10,800 across the population (compare with €310,000 for real estate wealth, 28 times as much); this rose to €85,000 for those households with any holdings of risky financial assets - however, these households at the same time had an average of over €650,000 in real estate assets, i.e. over 7.5 times as much.

Risky assets are therefore not a very important component of wealth for the majority of households; is it however realistic for households in the model to only be able to invest in deposits or real estate? This question is most important for those who already own property and need to decide what to do with extra savings. In the data, out of the households who own their main residence, 26.2% have other real estate property (average of €123,000) but only 9.2% (6.4% of the overall population) own risky financial assets but no other real estate property (average of €58,000), showing again that these households are also a minority in real life and excluding them does not significantly impact the wealth distribution.

2.4. RESULTS

To estimate the impact of various housing policies on outcomes such as home-ownership rates, wealth inequality and welfare (and its distribution), we rely on a baseline model as basis for comparison, and we then change some targeted parameters or features one by one, or in combination. We are interested in estimating the impact of the following housing policies:


- Tightening of borrower-based macroprudential regulation (2.4.2);
- Introduction of institutional investors (2.4.3), and interaction with macroprudential regulation (2.4.3);
- Changes in tax rates on rental income (2.4.4);

- Measures on the construction sector (2.4.5);
- A negative scenario (2.4.6)
- Changes in interest rates (C)

However, we first make a few general observations about the joint housing and rental markets, and inequality.

Housing demand is generally a decreasing function of house prices (if housing is a normal good), while it is increasing in rental return (which makes owning a house vs renting more attractive financially). The opposite holds for rental demand. Accordingly, if we ignore for now housing supply (suppose it is equal to N at all times), there can be a continuum of combinations of p_h and r_h that lead to an equilibrium in the housing market: e.g., a higher p_h , which decreases demand for housing, can be compensated for by a higher r_h to discourage renting. This would result in multiple equilibria as pictured in Figure 2.1.


Figure 2.1: *Illustration of multiple equilibria in the housing market when supply is irrelevant.*



This is confirmed by model simulation results, where housing supply was set exogenously to be equal to N_t at all times, while all other parameters are calibrated to the baseline specification. Figure 2.2 displays the level of housing demand for each p_h and r_h combination, confirming that demand increases with the rental yield and decreases with the house price. Here we can see how demand for housing reaches equilibrium at a number of combinations of p_h and r_h , ranging from low-price low-yield equilibria to high-price high-yield ones. Note how Figure 2.1 is simply a representation

of a slice of Figure 2.2, taken along the Z axis at the point where demand is 100% of the population.

Figure 2.2: *Surface area representation of demand for housing as a percentage of population under the baseline model, by p_h and r_h .*




Out of all these possible equilibria, are some preferable to others from a social perspective? It turns out that a low-price, low-yield equilibrium is strictly better than a high-price, high-yield equilibrium from the point of view of equality. In fact, when prices and yields are high, landlord households enjoy high levels of housing wealth from their properties and high levels of income from rent; meanwhile renters, who are poorer, find it harder to access home-ownership due to the high house prices, while also paying high rents. On the other hand in a low-price, low-yield world, renters would benefit from lower rents and easier access to housing, while wealthier households would have lower housing wealth and rental income; this would lower several measures of inequality. Indeed, we will see how in general any policy change that would lead to an equilibrium increase in p_h when keeping r_h constant, or vice-versa, will tend to lead to an increase in inequality, the opposite also being true. Housing not being a productive asset also means that no benefit accrues to the economy from more wealth being held in housing.¹³

To illustrate this relationship between the equilibrium levels of prices and yields, and inequality, we have taken a set of combinations of p_h and r_h

¹³This is a general remark rather than a finding of our model: remember that there is no capital in this economy.

that lead to an equilibrium (housing demand equal to 100% of the population as per Figure 2.2), and calculated a number of Gini coefficients for households in these economies. As Figure 2.3 displays, Gini coefficients rise (net of fluctuations due to randomness in the different simulations) together with the equilibrium house prices (which also correspond to higher rental yields), indicating increasing inequality. Consumption and income inequality increase significantly more steeply than wealth inequality, which however remains the highest throughout. This can be explained by lower-income households having higher marginal propensities to consume, as is commonly found in the literature, meaning that higher rents (reducing their disposable incomes) will in great part translate to decreases in consumption rather than in saving and therefore wealth.

Figure 2.3: *Gini coefficients for wealth, disposable income and consumption, for each (p_h, rh) combination leading to equilibrium ($H_d = N_t$) in the baseline model.*




2.4.1. The baseline model

We now return to the complete solution to the model. Once constraints originating from the construction sector (where p_h is a determinant of construction output), are taken into account, there is a single price p_h^* that makes sure in the steady state there is exactly one housing unit for each household. Figure 2.4 illustrates this. While the equilibrium price of a unit of housing is uniquely determined by the construction sector, the intersection with housing demand determines the equilibrium rental return r_h^* . This can be found, for instance, by intersecting the surface of Figure 2.2 with a $H_s = N_t$ plane corresponding to $p_h = p_h^*$ and finding the unique level

of r_h in that intersection that also satisfies $H_d = N_t$. With the calibration of our baseline specification, p_h^* is set to 5.8 times the aggregate wage, with the corresponding optimal rental yield being around 3.15%.

Figure 2.4: *The construction sector determines the optimal price, while intersection with housing demand determines the equilibrium rental rate.*



Note how, from these levels, an increase in housing demand by households (for example, due to an increase in preference for home-ownership) would shift the plane in Figure 2.2 up, or the $H_d = N_t$ curve in Figure 2.4 to the left, and lead to a lower equilibrium rental return, and vice-versa.


2.4.2. Tightening of macroprudential measures

The degree of tightness of the macroprudential measures determines, as well as the overall level of credit in the economy, whether some individual households are able to avail of mortgage debt to purchase housing, and if so how much debt they are allowed to take. This can have important distributional and welfare consequences, especially for those households who are close to, or just about meeting the LTI and LTV thresholds.

When macroprudential rules are tightened, some of the owner-occupier households who met the LTV and LTI criteria under looser regulation will no longer be able to access mortgage debt. These households, who previously contributed to housing demand, will now be part of the rental market. Even some landlord households may not satisfy the criteria for borrowing to finance the purchase of an extra housing unit, reducing their demand for extra housing (i.e. the rental supply) while at the same time the demand for rental accommodation increases.

This decrease in household demand, as outlined above, will shift the $H_d = N_t$ curve in Figure 2.4 to the right, resulting in a higher rental yield. The (forced) increase in the number of renters compared to a situation with looser thresholds leads to an increase in rental demand, pushing up rental prices. At the same time, the remaining existing landlords need to collectively own an additional number of housing units equal to the extra renters. In order to incentivise them to own more houses, the rental return needs to adjust accordingly. Figure 2.5 illustrates this graphically.

Figure 2.5: *Effect of a tightening of macroprudential policy forcing some households into the rental market. Seen in two different ways.*



Note how, in the absence of constraints arising from the construction sector, the shift in demand might lead to a decrease in price, an increase in yield, or a combination of both. Some empirical studies have found that indeed, in the short term borrower-based macroprudential policies lead to a decline

in house price growth, and indeed these types of macroprudential policies have in most cases been introduced in order to reduce the risk of a credit-induced unsustainable boom in house prices. While this might hold true in the short and medium run (these policies aim at reducing fluctuations within this time frame), our framework suggests that in the long run, the necessity of housing construction keeping pace with population growth means that prices essentially need to be anchored to a certain level. As a result, the long-run impact of a tightening in macroprudential policies is an increase in rental yields and no change in prices; i.e., higher rental costs in absolute terms.

With respect to the welfare and inequality implications of a tightening of macroprudential thresholds, the main channels through which different households are impacted by such policy changes are the following:

- The most directly impacted households, i.e. those who would have obtained a mortgage with looser policy but are instead forced to rent, lose in both welfare and monetary terms (they lose the additional utility from owning their house, and the rent is typically higher than equivalent mortgage repayments in the looser policy scenario).
- Households with enough liquid assets are generally not impacted by changes in the rules, while being benefited by an increase in rental yields if they own more than one house.
- Households that would have rented in any case (poorer) see their rental payments increasing.
- Generally, inequality is increased while welfare decreases for poorer households and increases for wealthier ones. Some households comparatively lose out when younger but are eventually benefited when older, while the negative consequences of the different policies are permanent for some others, such as lifetime renters.

These facts can be seen in Table B.1, which summarises the effects of lowering the loan-to-income threshold from the original 6 to a significantly tighter 3.5.¹⁴ In this simulation, the decline in housing demand generated by the tightening in the LTI limit leads to an increase in the equilibrium

¹⁴Recall that the equilibrium price to aggregate wage ratio - which does not change with tighter macroprudential rules as the construction sector remains the same - was 5.8 in the baseline specification. A 3.5 LTI threshold would thus make a mortgage inaccessible for a large proportion of the working population unless they have sufficiently large liquid savings.

rental yield, from the original 3.15% to 3.45%, while p_h remains 5.8 times the aggregate wage.

The distributional effects of the change in policy are most obvious from the allocation of housing: while under the 6 LTI scenario (the baseline), homeownership is 70%, it falls to 66.1% under the 3.5 LTI scenario, i.e. about 4% of the population become renters but would have been homeowners if less restrictive regulation had allowed them to borrow. The increase in the number of renters is matched by a corresponding increase in the number of houses owned by landlord households, who have been incentivised to own additional housing thanks to the higher rental yield.

There is, under the 3.5 LTI scenario, at the same time a higher percentage of the population paying rent, and also a higher level of rent paid by each renter. As a result of these dynamics, households at the lower end of the wealth distribution are the most impacted financially,¹⁵ and see both their wealth and their consumption decline.

Households who own one housing unit, typically in the middle of the wealth distribution, are relatively unaffected by the changes in policy as they neither pay rent nor receive rental income.¹⁶ This suggests that, for households who do end up owning a house eventually, the tightening in the LTI threshold merely represents a shift in the age at which they are able to purchase (with corresponding decrease in consumption and utility in the additional years of renting). In fact, they may overall benefit from the change in policy should they eventually become landlords. However, there are a number of households who, as a result of the tightening, become lifetime renters.

At the upper end of the wealth distribution, there are more households owning 2 or more housing units, and they receive a higher rental yield on each of them compared to a situation under an LTI threshold of 6. This leads to a significant increase in both their wealth and their consumption (and thus utility). All the above point to a marked increase in inequality as a result of the change in policy. Indeed, the Gini coefficients on housing wealth, total wealth, disposable income (after rent payments), and consumption all increase considerably. Utility levels decrease sharply for the

¹⁵Renters always tend to be the poorest households, as the extra utility from homeownership means that it is typically optimal for households to own a house if they can afford it.

¹⁶As previous renters they had been affected earlier in their life, but this does not seem to translate into significant deviations in wealth or consumption relative to comparable households under the 6 LTI simulation.

poorer and increase only slightly for the wealthier.

What happens when the macroprudential limit being tightened is instead the loan-to-value ratio? While the LTI threshold limits some households' access to credit due to their income levels, the LTV threshold does so based on their liquid wealth. Some young, high-income households, for instance, may be able to meet the LTI requirements but not the LTV ones as they have not yet been able to accumulate sufficient liquid wealth. The need for a larger deposit both delays home-ownership and incentivises saving behaviour. On the other hand however, most other households would typically be able to meet a tighter LTV limit more easily than a tight LTI limit. Consider the following: at prices at a multiple of 5.8 times the average wage as per baseline specification, an average-wage household saving 15% of their income every year would need about 4 years to save for a 10% deposit and 8 for a 20% deposit on a house at such a price. Tightening the LTV requirement could therefore delay their purchase. However, with a deposit of 10% they would need a loan for 90% of the house price, i.e. about 5.3 times their income, while with a 20% deposit the required loan would be about 4.7 times their income. Both would be inaccessible with a tight LTI limit in any case. Indeed, for many households, a tight LTI limit means that, in order to be able to get a loan, a significant part of the purchase price of the house would need to come from a deposit, and this is typically higher than the deposit required by the LTV limit itself.¹⁷ Results of lowering the LTV limit from 90% to 75% (in Table B.1) confirm that such a policy change does not have a particularly large effect on inequality; however, it is likely that tightening LTV limits further would start having a significant impact at a certain point.

2.4.3. Institutional investors

As seen, the main features of institutional investors in our framework are that (i) they are externally funded and deep-pocketed; (ii) they are targeting a minimum return r_h^{ii} (which depends on their cost of funding, r^i , and other factors as detailed in equation 2.16), at or above which they will acquire as many units as possible; (iii) they will do so regardless of price (as there are no short-term price fluctuations within our model).

¹⁷In the same example of the average-wage household, an LTI limit of 3.5 with a purchase price of 5.8 would require a deposit of 2.3 times the household's income, or about 40% of the purchase price, i.e. the LTI limit itself effectively enforces an LTV limit of 60% on this particular household.

Due to these characteristics, the presence of institutional investors will provide an anchor to the rental yield, in any situation in which this would naturally be higher than r_h^{ii} in their absence. Suppose for instance that in the absence of institutional investors, demand for housing would be insufficient at r_h^{ii} , meaning that not enough households would be willing to invest in rental property for that rental return, leading to a shortage of housing demand and excess rental demand.

While without institutional investors this would lead to r_h adjusting upwards, the institutional investors are happy with the return and will therefore purchase any units that are not claimed by any households, essentially anchoring the rental yield at r_h^{ii} . This would be the highest rental yield achievable, with any $r_h < r_h^{ii}$ also possible and entailing no presence of institutional investors in equilibrium. Figure 2.6 displays the mechanism at play.

In a situation in which institutional investors are present, in combination with constraints from the construction sector, the marginal household would essentially be facing a given house price as well as rental yield in the long term. Households' welfare and inequality will mainly be affected in the following ways:

- Renters, especially those who would have never owned a house throughout their life (i.e. the poorest), benefit from a lower rent.
- On the other hand, household landlords receive a lower return on their investments and have generally lower housing wealth (fewer houses for investment).
- There might also be a number of households who now rent and would have been homeowners otherwise. Homeownership levels decrease.

As these changes effectively lead to an increase in disposable income for renters, while also lowering investment income and housing wealth for landlords, they lead to significantly lower income and consumption inequality. The extent to which this happens depends on the distance between the equilibrium rental return without institutional investors, and r_h^{ii} : the greater this is, the larger the effect.

Figure 2.7 displays how the distributions of housing, consumption and wealth change when we add institutional investors to the model. We assume r^{ii} equal to 1/3%, leading to $r_h^{ii} = 2.4\%$ as per equation 2.16. As

Figure 2.6: The presence of institutional investors pins the rental return. This might reduce household sector home-ownership and wealth levels at the top, but also rents.





Figure 2.7: *Impact of the presence of institutional investors in the housing market on the distribution of housing wealth, consumption and total wealth, with LTI of 6 and r^H of 1/3%.*



this is significantly far from the 3.15% rental yield that would prevail in the absence of institutional investors (see above), the effects are quite noticeable. The entrance of institutional investors into the housing market leads to fewer households owning two or more houses, and more of them owning 1 or none. Homeownership indeed declines by 3 percentage points compared to baseline. Wealth declines across the whole distribution (as less housing is in the hands of households, and its return is now lower), but more so for the wealthier households. Despite the fall in housing and overall wealth, a lot of households gain in terms of consumption, with an average household in the bottom 30% of the consumption distribution consuming an extra 5% of the aggregate wage compared to the same type of household under the baseline model. This is thanks to significantly cheaper rents when institutional investors are introduced. On the other hand, the wealthiest households significantly reduce their consumption, as their wealth suffers a large decline.

Overall, the presence of institutional investors, bringing down the equilibrium rental yield, can lead to a decrease in income and consumption inequality thanks to lower rents, at the cost of reduced homeownership, wealth and domestic ownership of the housing stock. Relaxing the strong assumption that they are only allowed to purchase those housing units that are not demanded by households could however significantly change

things, and generally for the worse.¹⁸ In this case, in fact, homeownership by the domestic households would fall even further, and for some households that would have not been the optimal choice in terms of their utility.¹⁹ A further complication arises if, for whatever reason, the institutional investors do not face competition: in this case they would not force the rental rate down (no need to compete with other similar investors to offer renters a lower rate), but they would still purchase as many housing units as possible, thus decreasing the homeownership rate as well as domestic wealth, but with no benefit for domestic renters.

Note also that while we call the investors in this scenario ‘institutional investors’, they might also represent the Government buying social housing and renting it out at a favourable rate, or at cost. In this case, the resulting equilibrium rental rate might be even lower than r_h^{ii} , since the Government is not looking for a return, and does not need to pay taxes to itself; that would lead to even less inequality, and in addition the housing wealth, as well as all rental income, would stay within the domestic economy.²⁰ These are all scenarios worth exploring in further work.

With tightening of macroprudential rules

When the presence of institutional investors is coupled with a tightening of macroprudential regulation, the effects of policy tightening can be significantly different than when demand for housing is only determined by households. We saw above how when macroprudential policy is tightened, the equilibrium rental yield increases due to additional demand for rental and less demand for housing at each price.

However, the presence of institutional investors (before tightening) implies that the equilibrium rental yield is r_h^{ii} . Lower macroprudential thresholds require an increased rental yield for the housing market to clear, as more household landlord demand is required and this can only be incentivised with higher returns. With institutional investors this is not necessary, however: they will purchase any extra units that are no longer demanded

¹⁸Suppose for instance that who ends up buying a property is decided by a lottery if there is excessive demand: in this case, institutional investors would be in for a chance to get every single housing unit. We do not simulate this scenario for brevity.


¹⁹That is instead the case if institutional investors are restricted as detailed above.

²⁰In this situation, two separate rental markets might develop, with two different rental rates, unless social housing is open to every household equally and allocation between the public and the private rental markets is random. If rental markets become segregated, this could allow the very poorest to rent at lower prices while landlord households maintain their returns in the private rental market.

by households and maintain the yield at r_h^{ii} .

In this case, the net result of a tightening in macroprudential policy would then be, differently from the household-only version of the model, a drop in the number of household landlords (as well as homeownership in general) but no change to prices, as illustrated in Figure 2.8. The households who would have been able to get a mortgage with looser regulation (and would have decided to do so) lose the extra welfare from home-ownership. Also, a number of landlords who would have got extra properties using mortgages under looser regulation lose rental income from these properties; this income goes instead to the institutional investors. Note as well that the rental income paid to institutional investors leaves the domestic economy and therefore does not finance any other household's consumption or saving.²¹

Figure 2.8: A tightening of macroprudential policies when institutional investors are present lowers the number of housing units owned by domestic households (whether they are an investment or owner-occupied).



Due to the low rental yield equilibrium imposed by the institutional investors, household demand for housing is significantly lower compared to baseline, as r_h^{ii} provides little prospective return to landlords, while many renter households find renting more convenient. The tightening of macroprudential regulation further reduces housing demand, as we have seen in Section 2.4.2. This results in a sharp drop in homeownership, which falls from 67% in Section 2.4.3 (and 70% in the baseline) to only 61.9% (compare Table B.1). At the same time, however, measures of inequality

²¹The same holds for the previous section. However, recall that this does not impact production as there is no capital in the economy.

(apart from housing wealth inequality), barely move. This is in contrast to the strong increase inequality between the baseline and the scenario in Section 2.4.2.

The difference from Section 2.4.2 in terms of effects of macroprudential policy on inequality is that those who are pushed towards renting due to tighter policy do not have to pay higher rents, as was instead the case when the LTI limit was tightened from the baseline simulation. Similarly, a prospective landlord who is unable to access a loan loses less under this scenario, since prospective returns are lower anyway. In other words, the welfare losses from not being able to access a loan are much lower now that the opportunity cost of not being able to own a house (i.e. the rental yield) is also lower. As a result, we have that tightening macroprudential regulation with institutional investors already present has a much smaller effect on inequality; conversely, introducing institutional investors when macroprudential policies are already tight greatly reduces income and consumption inequality, but also homeownership, with an unclear effect on wealth and housing wealth inequality.²²

2.4.4. Tax on rental income

In our baseline specification, we have rental income tax standing at 40%. This creates a large wedge between what renters pay, and the income eventually accruing to landlords. Reducing it might potentially reduce the burden to renters, increase the net return to landlords, or both. Here, we look at the effects of changing this tax rate within our framework, with the caveat that the high level of complexity of the (fiscal) implications of such a tax would warrant further in-depth study.


Does the majority of the tax burden fall on renters, increasing their payments, or does it fall mostly on landlords, by reducing their effective after-tax rental yield? Who bears the burden of the tax can have potentially very different implications for inequality, and it can depend heavily on calibration and indeed the interaction with other policies. In addition, in the real world the revenue generated from such a tax can be used in a number of ways that can further contribute to or mitigate inequality. For instance, the Government could use the tax revenue to decrease other tax rates, subsidise house building, or on social spending (including social housing), etc., all with very different implications, some of which we

²²This depends on the degree of tightening of policy and the level of r^{II} .

separately discuss in other sections.

The interaction with other policies can be significant. Are expenses on maintenance, property taxes, or mortgage income deducted from the tax? Do institutional investors pay the same tax rate on their rental income or do they pay a (higher or lower) tax on their profits? For the sake of this exercise, we limit ourselves to observing what changes when the tax rate on rental income is reduced, within the setting of our baseline model, from 40% to 20%. In order to concentrate on the change in the tax rate alone rather than the use of the revenue it generates, we assume that any shortfall in tax revenue from the reduction in the tax is taken from the provision of public services, which as explained above do not impact household decisions, or their income and wealth.

Figure 2.9: *Lowering τ_{rh} benefits both renters and landlords, but not in equal measure*



The result of the reduction in τ_{rh} is, unsurprisingly, a decrease in the equilibrium rental yield. Both renters and landlords gain thanks to the reduction in the tax,²³ with the first paying a lower rent and the second receiving a higher after-tax rental income relative to the baseline specification. From the original 3.15%, the equilibrium (gross) rental yield falls to 2.60%. This represents a fall of over 17% in the level of rental payments by renter households compared to baseline, with rents as a proportion of average income falling from 18.3% to 15.1%, meaning that each single

²³Again, we are agnostic as to what the Government does with the tax revenue. In the real world it could as well use the revenue in ways that benefit all households more than they are harmed by the tax.

renter is better off by about 3% of average income (see also Table B.1 for more statistics). Despite the decrease in yields, landlords are also better off because thanks to the decrease in tax, their after-tax return is still higher compared to the after-tax return under the baseline simulation. From getting taxed 40% on a yield of 3.15% (net yield, 1.89%), they are now taxed 20% on a yield of 2.60%, with net yield of 2.08%. That is a 10% increase in after-tax income.

All households (apart from owner-occupiers) benefit from a reduction in the tax, but renters significantly more so than landlords under this calibration. We show this in Figure 2.9, where we have evaluated the equilibrium rental yields and net rental yields (i.e., respectively, what renters pay and what landlords receive after tax) at income tax rates τ_{rh} of 40% and 20%, across a number of different levels of p_h . At all p_h levels, the gains made by renters are several times as large as those made by landlords when the tax rate is decreased. This suggests that the burden of the tax falls mostly on renters. It thus appears as though, despite the fact that it is a tax levied on landlords' income, rental income tax (at least under our calibration) is a regressive tax, as it hits poorer households more than wealthier ones (in monetary terms, but even more so in proportion to their income). Indeed, while housing and wealth inequality increase with the reduction in the tax (presumably, because landlords have an even higher net return now),²⁴ income and consumption inequality both significantly decline, and utility rises on average but especially so that of poorer households.


2.4.5. Measures on the construction sector

We have seen above how the construction sector determines the steady state level of p_h regardless of housing demand. The specific level of p_h that allows for housing supply to match the population therefore must be dependent on characteristics of the production function of the construction sector. As low-price, low-yield equilibria are preferable from an equality standpoint, inequality can be reduced if construction expands at any given price level. When this happens, a lower price is required for the construction sector to be able to meet housing supply requirements, leading to a lower equilibrium price. Figure 2.10 displays this graphically.

An increase in construction output can be achieved in a number of different ways, with similar implications. We look at a couple of them below.

²⁴And also, there is effectively more money going around, because we did not change any other tax rate.

Figure 2.10: Expanding construction permits leads to a lower p_h allowing for $H_s = N_t$. This in turn leads to a lower equilibrium r_h



In our first simulation, we assume that the number of land permits is increased by 20% in every period. This leads, as predicted, to a decrease in the equilibrium p_h and r_h , from 5.8 and 3.15% to 5.3 and 2.95%. As suggested from Figure 2.3 and displayed in Table B.1 in the Appendix, moving to a lower-price, lower yield equilibrium significantly lowers Income and consumption inequality. This is because lower house prices coupled with lower rental yield leads to significantly lower rental payments (which are $p_h \times r_h$), which benefits poorer households who can enjoy higher disposable income and consumption. Rent as a percentage of the average wage falls from 18.3% to 15.6%.

As it is cheaper, households can more easily afford housing (if they so wish), with or without loans, but note that this does not necessarily lead to higher homeownership: some households are happy to rent at the lower r_h , as the *financial* incentive to own housing falls significantly. Utility improvements compared to baseline are indeed largest for renting households, and the homeownership rate actually falls. Due to the lower price and return of housing (which still has a higher return than deposits), the wealth of homeowners and landlord households is negatively affected by the lower house price and lower returns from the properties they rent; however, consumption only decreases at the top 20% of the distribution, while it increases for everyone else.

Overall, the decrease in house prices seems to benefit more households than it harms, and those households who are worse-off compared to when prices are higher remain among the wealthiest in any case. For instance,

the average increase in consumption for households in the bottom 25 consumption percentiles is of about 1.5% of the annual aggregate wage (2% for the lowest 10 percentiles) relative to households in the same percentiles under the scenario with lower permits.

Next, we simulate an increase in the productivity of the construction sector, with A_h going from 1.3 to 1.4. While this is not per se a policy, there are a number of real-life Government policies that can help achieve an improvement in construction productivity, for example providing efficient planning and legal systems in which the construction firms can operate. The outcome of this scenario is very similar to an increase in land permits, as seen from Table B.1: indeed, to households it obviously does not matter how the expansion in supply is achieved. As above, who benefits from the lower r_h and p_h are renters (due to lower rental costs), as well as homeowners with only one property (if housing is cheaper, they can buy earlier and enjoy the additional utility from housing), while landlord households lose on wealth and consumption.

On this simulation, we further tighten LTI limits to 3.5 to see how the policies interact. As happens for the baseline simulation, even with the increased housing supply at all p_h levels, tightening the LTI threshold from 6 to 3.5 will lead to lower demand for housing, and an equilibrium with higher rental yield (namely 3.25%), the same house price (5.2) and higher inequality. However, rent per average wage remains well below baseline (16.9% vs 18.3%). Housing, wealth, and income inequality all increase compared to baseline, but consumption inequality decreases and utility improves for the poorest.

Because house prices in equilibrium are lower thanks to the increased productivity of the construction sector, fewer households are affected by the LTI limit even when it is tightened: with p_h of 5.8 and LTV of 90%, a household needs an income of 0.87 for a 90% loan at 6 LTI and 1.49 for an LTI of 3.5; with p_h of 5.20, the equivalent incomes required fall to 0.78 and 1.34, respectively.

Indeed, in this particular case the increase in construction sector productivity seems to more or less offset the effect of tightening the LTI limit, leading to a situation that is comparable to the baseline in terms of inequality (and indeed a slightly better outcome in term of utility, but with a lower homeownership rate (compare Table B.1)).

Other policies that can achieve similar results to an expansion in land

permits or increases in productivity, by affecting the construction sector rather than demand, include, for instance, subsidies on house prices (this however generates a cost to the Government, and therefore a trade-off), or a decrease in the cost of materials or labour in the construction sector.²⁵

2.4.6. A not implausible negative scenario

In the simulations above, we have analysed the distributional impact of policy changes around a relatively favourable baseline scenario. The reality is, however, that a number of advanced economies are far from our baseline scenario at the moment, being concurrently faced with insufficient supply of housing, high rents, tight regulation, high taxation, etc. Here, we simulate a scenario where we combine a number of the above policies in an unfavourable way, to show how quickly they can add up and contribute to significant inequality. Namely, we set land permits to 80% of the baseline rate, we allow for institutional investors (but with a relatively higher r^i of 1%), we set the LTI limit to 3.5, and we increase the rental income tax rate to 45%. This combination is unfortunately far from implausible.²⁶

Under this scenario, a large number of households end up being much worse off. As displayed in Table B.1, all inequality measures increase substantially compared to the baseline and all other scenarios examined, with the homeownership rate also falling from 70% to 63%. The equilibrium p_h jumps to 6.55 and the equilibrium rental yield reaches 3.82%; together, this implies that renters spend 25% of the average wage on rent under this scenario.²⁷ For all homeowner households, wealth increases substantially compared to the baseline simulation, as housing is more expensive but also of higher return - in fact, the wealthiest households gain the most wealth, and increase their consumption levels accordingly. On the other hand, renters see their limited wealth further reduced due to the high rental rates. Their net income and consumption - and thus utility - all fall sharply.

What can policymakers do when faced with a similar scenario of high

²⁵While labour is not an input for construction firms in this model, a reduction in its cost would of course improve supply in a real economy.

²⁶Indeed, it is not far from the current situation in Ireland: household landlords pay a marginal tax of around 50% on rental income when including all types of taxation (but mortgage repayments are deductible); the planning process is slow and inefficient; construction sector productivity is low by euro area standards; a large proportion of the rental market is in the hands of foreign investment funds; LTI and LTV limits for first-time buyers are respectively 4x and 90% currently.

²⁷Note that renters typically earn much less than the average, so this is an ever heavier burden to them.

inequality, and a large number of vulnerable renters? We hope that results highlighted in this paper can shed some light into this. We believe that supply-side measures should be prioritised, as we have shown that when the construction sector can build more housing and/or at lower cost, outcomes for households are unequivocally and strongly positive. A review of rental income taxation (for household landlords) could end up benefiting renters and decreasing inequality. Appropriate regulation of institutional investors, especially if they operate within a small market with limited competition, might be warranted: while our results in 2.4.3 point out that they can have certain positive effects, it is important to avoid a situation in which they end up dominating the rental market. On the other hand, provision of public social housing would provide the same benefits we identified in institutional investors, with reduced welfare costs and risks.

2.5. CONCLUSIONS

In this paper, we have analysed the effect of a number of different housing market policies on inequality, employing a heterogeneous agent overlapping-generations life-cycle model with housing and rental markets. We find that through limiting availability of credit to lower income households, macroprudential regulation has a meaningful impact on rates of home ownership, with implications for income, wealth, and consumption inequality. Housing policy targeting supply, minimizing costs associated with new home building can work to improve outcomes for potential homeowners at the bottom of the distribution. The same cannot be said for demand-side policies which, in the face of limited supply response, will largely work to benefit the wealthy.

We further explore the effect of the introduction of institutional investors into the rental market, and find that they can decrease inequality by lowering equilibrium rental rates, at the cost of reduced homeownership. Taxation on rental income, targeting the incomes of the wealthiest, might actually harm renters the most by pushing up rental payments. Different policies can lead to different trade-offs between different types of inequality, as well as homeownership rates. An optimal policy mix can be achieved by combining these policies in an effective way; in this light, we note that supply policies aimed at improving the availability and reducing the cost of housing are essentially always positive.

We note that there are many unanswered questions remaining. In future research, we plan to expand our analysis to further experiment with different policy impacts, and to narrow down our focus on some of the policies we have already discussed here, but in more detail. The model setup outlined in this paper is very flexible and lends itself to further experimentation, with different calibrations and additional features; we plan to make use of this to explore further areas for investigation. For instance, while preliminary work suggests that help-to-buy schemes will not improve inequality, we plan to test this. Additionally we wish to better understand the degree to which various policies might work together.

Ultimately we emphasise that any study that wishes to understand the role of policy to affect inequality should take the housing market seriously. For better or worse it is an asset that makes up the majority of households' wealth and is therefore absolutely central to any serious attempt to understand and minimise such distortions.

3 Monetary policy shocks across the wealth distribution: insights from new European data

3.1. INTRODUCTION

Few developments have as broad an impact on all aspects of the economy as monetary policy does. Whether directly through interest rate movements, or indirectly through employment or inflation, every household is affected one way or another by monetary policy and its transmission across a number of different channels. However, depending on factors such as employment and income, as well as the amount and relative allocation of asset holdings, households can be impacted by (and respond to) monetary policy in very different ways, and systematically so.

What happens to households' wealth in response to monetary policy? There is a large and growing literature¹ investigating the heterogeneous effects of monetary policy on households, based on characteristics such as income level, wealth, asset holdings and indebtedness. Such heterogeneous responses to monetary policy naturally lead to developments in wealth inequality among households with different economic characteristics.

In spite of the abundance of studies and theoretical understanding of the heterogeneous effects of monetary policy on households' wealth, data limitations have made it quite difficult to gain insights into the distributional effects of monetary policy at the aggregate level.² In other words, not much is known about how large groups of households with similar characteristics

¹As well documented in [Colciago et al. \(2019\)](#).

²The World Inequality Database ([Alvaredo et al., 2024](#); [Saez and Zucman, 2016](#)) is a notable undertaking in this direction and a great source of global data on the distribution of income and wealth. However, it is not suited not suited to study the effects of monetary policy on inequality due to its low frequency.

(such as their level of wealth, which is the primary focus of this paper) respond to, or are affected by, monetary policy.³ Until recently, studies have been largely based on survey data, simulations and models. This paper contributes to the literature on monetary policy and inequality by providing empirical findings that represent a bridge between the current micro-based (such as [Casiraghi et al., 2018](#) and [Slacalek et al., 2020](#), who use a simulation approach) or aggregate empirical findings, and the results of models such as [Auclert \(2019\)](#) and [Luetticke \(2021\)](#).

To do this, we employ the Distributional Wealth Accounts recently developed and published by the European System of Central Banks. This dataset combines detailed (but infrequent) evidence on households' finances provided by the Household Finance and Consumption Survey with timely (but aggregate) data from the Quarterly Sector Accounts of the euro area. As a result, this new dataset provides us with quarterly time series of household wealth and asset holdings, broken down by a number of groupings (namely, wealth level, employment, and home-ownership status). The (relative) high frequency of this new data allows us to study the impact of monetary policy on the wealth and balance sheet composition of households across the whole wealth distribution, and to identify the effect of monetary policy on wealth inequality at an aggregate level using time-series techniques, which is a novel approach when it comes to distributional considerations in a European context.

Our approach is similar to the one used in [Franconi and Rella \(2023\)](#) for the United States.⁴ However, our results are novel for the euro area, and we also significantly expand the treatment of this topic by explicitly investigating the impact of monetary policy shocks on households' investment decisions, portfolio allocations, and by carrying out analyses by household homeownership and employment status. Additionally, we also test whether there is asymmetry in the distributional effects of contractionary versus expansionary monetary policy.

Our results are broadly in line with the rest of the literature and with the effects of monetary policy changes on the value of assets typically implied by economic theory.

An increase in policy rates is expected to cause a fall in the price of risky

³Unless these effects are simulated.

⁴The Distributional Financial Accounts, the US equivalent of the euro area Distributional Wealth Accounts, have been available for longer and thus comprise longer time series.

financial assets (e.g., equity) and house prices, thus lowering the wealth of households holding such assets.⁵ Equity prices are expected to fall as a result of contractionary monetary policy because of the increased return on alternative risk-free assets. If risk premiums stay the same, the future expected return on equities also needs to rise, and for this to be possible the asset price needs to decline accordingly. In addition, equity prices fall because restrictive monetary policy is expected to reduce most firms' expected future revenues due to a general slowdown in the economy (Maio, 2014). In the case of house prices, higher rates on mortgages discourage some potential mortgage buyers from the housing market, therefore reducing demand. At the same time, the differential between rental yields and risk-free returns narrows, thus making housing as an investment also less attractive.⁶ Finally, bond prices are inversely correlated to bond yields, which closely follow movements in policy rates.

With equity, housing and bond prices falling in response to restrictive monetary policy, we expect to see these negative asset price movements reflected (at least in part) on households' wealth, with different households being differently affected based on the size and weight of these assets on their balance sheets. Once we know what happens to asset prices, the distributional effects of such movements are well studied. For instance, Adam and Tzamourani (2016) show that an increase in house prices would tend to benefit households of all levels of wealth, but more so those in the middle of the wealth distribution. On the other hand, increases in equity and bond prices would overwhelmingly benefit those at the top, as these assets account for a larger share of richer households' wealth.

In line with the findings described above, our results indicate a fall in the value of most households' risky financial assets in response to a restrictive monetary policy shock. Despite being the most exposed to equity, however, we actually note an initial *increase* in the levels of risky assets for households in the top 10% (and less so for decile 9) of the wealth distribution. Consistently with theory, we also note a marked decline in non-risky fi-

⁵See, among others, Ehrmann and Fratzscher (2004) and Ioannidis and Ktononikas (2008) for what happens to equity prices, and Mishkin (2007), Jordá et al. (2015) and Koeniger et al. (2022) for how monetary monetary policy affects house prices and the housing market.

⁶The response of asset prices to monetary policy is largely predictable qualitatively, but highly uncertain quantitatively: as Paul (2020) highlights, the responses of house and equity prices to monetary policy can vary greatly across time and based on circumstances in the housing and stock markets themselves. Furthermore, there are wide ranges of estimated coefficients in the literature, both across and within economies - see Williams (2015), for instance.

financial assets across all households in response to a contractionary shock. Debt also falls substantially for all but the wealthiest households. For most groups, we find that housing wealth declines in response to a shock towards the middle of the horizon, however the effect (as expected, since housing is less liquid) is not as strong as for financial assets and is generally non-significant.

Apart from movements in asset prices, monetary policy shocks also affect the growth of credit and money circulating in the economy: consider the classic credit channel of monetary policy as per [Bernanke and Gertler \(1995\)](#); and [Drechsler et al. \(2017\)](#) for the behaviour of deposits. Indeed, this is one of the main channels of transmission of monetary policy. Consequently, we expect a contractionary monetary policy shock to negatively affect the levels of both debt and risk-free assets, this time via a lower volume of new loans and reduced flows into deposits rather than via price effects. The decline in credit eventually leads to general equilibrium effects, with output and inflation both falling ([Cloyne and Hürtgen, 2016](#); [Romer and Romer, 2004](#)). [Slacalek et al. \(2020\)](#) find that low-income households are more strongly affected by these general equilibrium effects as they are more likely to become unemployed as a result of a recession, and are also relatively more reliant on labour income.

Monetary policy can also affect household wealth through a number of other channels, starting from how it impacts the financial income and debt repayments of savers and borrowers (see [Auclert \(2019\)](#) and [Tzamourani \(2021\)](#)). Based on whether a household has higher holdings of interest-bearing assets (i.e. deposits and bonds) or liabilities (mortgages and other debt), their income⁷ would benefit or suffer from an increase in interest rates. In the former case, increased interest income from a household's deposits would more than offset any increase in debt repayments, while the opposite would be true if (variable-rate) debt levels are larger than deposits. As savers are typically wealthier, increases in policy rates, which benefit them and hurt borrowers, increase inequality. Monetary policy's effect on inflation has distributional implications, too: since the values of deposits and debt are both nominal, an increase in prices leads to a fall in the real value of both. Inflation is therefore good news for borrowers and bad for (deposit) savers, thus reducing inequality through this channel,

⁷And therefore wealth: higher income translates into higher savings, wherever the marginal propensity to consume is less than one, which is almost always the case ([Carroll et al., 2017](#)).

as borrowers tend to be less wealthy.⁸ Therefore, a restrictive monetary policy shock - which also delivers a negative unexpected inflation shock - would conversely hurt borrowers and increase inequality through this channel.

Monetary policy further leads to changes in households' saving and investment behaviour. Households in different wealth groups might indeed react differently to changes in asset prices and interest rates depending on factors such as their income, employment status, the relative importance of these assets in their portfolios, and the availability of liquid assets. For instance, a typical 'non hand-to-mouth' household (i.e., a wealthy household with significant liquid assets) would be able to respond more optimally to shocks than a household with more liquidity constraints, whether wealthy or not.⁹ In general, a rise in real rates would encourage savings via a direct intertemporal substitution effect. General-equilibrium income effects, however, are found to be stronger than the direct intertemporal substitution and interest exposure effects for hand-to-mouth households in [Ampudia et al. \(2018\)](#): while the real return on saving has increased, these households cannot take advantage of this; in addition, their incomes actually decline in response to contractionary policy due to equilibrium effects, leading to actually less saving and consumption.

Unconventional monetary policy can have substantially different repercussions on households compared to conventional interest rate monetary policy, as it works through very different channels. [Casiraghi et al. \(2018\)](#) find a U-shaped effect of quantitative easing on wealth across the distribution, with households at the bottom benefiting because of increased income and reduced cost of debt, while those at the top benefit from capital gains on financial assets. Overall, the effects on wealth inequality are found to be negligible. Through a micro-simulation on euro area survey data, [Lenza and Slacalek \(2024\)](#) find that quantitative easing policies reduce income inequality through increases in employment at the bottom of the income distribution, but they do not find statistically significant changes in the Gini coefficient of wealth despite the different portfolio holdings of households at different levels of the wealth distribution.

The rest of the paper is structured as follows: Section [3.2](#) describes the DWA

⁸See [Doepke and Schneider \(2006\)](#) and [Ferreira et al. \(2024\)](#). As per the latter study, inflation also affects wealth inequality through income and consumption channels.

⁹As well modelled in the HANK literature. See, for instance, [Kaplan et al. \(2018\)](#), [Luetticke \(2021\)](#), and [Slacalek et al. \(2020\)](#).

dataset and provides an overview of the composition and distribution of wealth in the euro area. Section 3.3 discusses the empirical strategy applied in analysis, and Section 3.4 details the corresponding results. Finally, Section 3.5 concludes.

3.2. DATA

Our primary source of data is the Distributional Wealth Accounts ('DWA') experimental dataset, published for the first time by the European Central Bank in January 2024. The DWA provides quarterly data on the amount of assets and liabilities held by households in euro area countries (excluding Croatia), as well as on their distribution across different sections of the population. Households are broken down according to three dimensions:

- **Net wealth deciles:** bottom 50% altogether, deciles six, seven, eight, nine, and the top 10% of the net wealth (i.e., total assets minus total liabilities) distribution;
- **Tenancy status:** Home owners (including part-owners) and tenants;
- **Working status:** employees, retirees, self-employed, unemployed and other.

Importantly, DWA series are constructed to be consistent with the existing macroeconomic aggregates for the household sector as a whole. In fact, the the dataset's methodology establishes a conceptual bridge to match the wealth instruments (assets, both financial and non-financial, and liabilities) reported in official macroeconomic statistics with items covered by specific surveys of household finances in each country. Information on the distribution of those comparable wealth instruments across households is obtained from the Household Finance and Consumption Survey¹⁰ ('HFCS') and then reflected onto Quarterly Sector Accounts ('QSA') series for the aggregate household sector which, otherwise, would not give any insights on distributional aspects.¹¹

As a result, the DWA provides, for the first time, a timely and granular view of the distribution of wealth in the euro area, consistent with the

¹⁰For more information on the methodology adopted in the latest wave of the survey, see ([Household Finance and Consumption Network, 2022](#)).

¹¹The construction process of the DWA and the assumptions adopted are described in greater detail in the 'DWA Methodological Note' published on the European Central Bank's Data Portal (together with an 'Overview Note' and a 'Frequently Asked Questions' document).

existing QSA statistics. In addition to time series for the stock of assets and liabilities, the dataset also includes some derived indicators such as the Gini index of wealth inequality, per capita amounts, and debt-to asset ratios by household group.

The DWA has been introduced in [Blatnik et al. \(2024\)](#) for the euro area, and its results have been described in greater detail in specific publications for the individual countries in the panel (e.g., [Kavonius and Topo, 2024](#); [Mesnard, 2024](#); [Moreno, 2024](#); [Neri et al., 2024](#)).

Table 3.1 displays a typical household’s balance sheet composition following the DWA concept of wealth. Some items that are otherwise covered in national accounts (such as pension entitlements or cash holdings) are not captured in the dataset due to lack of survey data. As such, they are also excluded from our analysis.

Table 3.1: Breakdown of households net wealth items in the DWA

Assets (A)	Liabilities (B)
<i>Non-financial</i>	<i>Debt</i>
Housing wealth	Outstanding mortgages
	Outstanding non-mortgage debt
<i>Risk-free financial</i>	
Deposits	
Debt securities	
<i>Risky financial</i>	
Listed shares	
Investment fund shares	
Life insurance and annuity entitlements	
<i>Business wealth</i>	
Financial business wealth	
Non-financial business wealth	Net Wealth (A-B)

We classify the ten DWA items into five broader wealth ‘types’:

- **Non-financial:** consisting of residential housing assets;
- **Risk-free financial:** consisting of financial assets with very low risk (or none at all) with predictable returns agreed in advance¹²;

¹²For classification purpose only, we consider debt instruments (e.g., bonds) to be held until maturity, and not to generate returns due to valuation changes. This assumption appears consistent with the view of households using debt instruments as ‘safe’ vehicles for saving.

- **Risky financial:** these assets generate returns for the owners primarily through revaluations and nominal changes in their value. They generally yield higher return rates, but also exhibit a greater degree of volatility;
- **Business wealth:** this category includes equity ownership in unlisted businesses as well as non-financial assets owned by households and used for production purposes;
- **Debt:** consisting of mortgage and non-mortgage loans.

3.2.1. Composition, distribution and dynamics of wealth in the euro area


The DWA gives important insights into how wealth is distributed in the euro area. These facts (neither new nor surprising to those familiar with results from the HFCS, in spite of the additional time series component given for the first time by the DWA) go a long way towards explaining and rationalising some of the results we will be showing below.

Figure 3.1 clearly highlights the absolute dominance of housing assets in households' portfolios at all levels of wealth, but especially so in the 'middle 40%' wealth group (i.e., deciles 6 to 9), broadly representing the middle and upper-middle class. For these, housing represents about three quarters of total assets. Housing is also by far the most important asset for households in the bottom half and the top 10% of the wealth distribution. The dominant role of housing is a feature in every single country in the euro area.

We also note how the relative importance of debt decreases as households get wealthier, with the bottom half of the wealth distribution having by far the highest level of debt as a share of assets. Because this consists almost entirely of mortgage debt (compare Figure E.1 in the Appendix), it follows that wealthier households typically own housing outright, while those in the bottom 50% own housing with a mortgage - and indeed have only very limited housing equity.

The composition of assets is strikingly similar across all households apart from those in the top 10% of the wealth distribution. Indeed, on average, households' assets for the bottom 90% can be characterized as being three quarters housing and one quarter deposits, with risky assets and business wealth playing only a very marginal role (albeit increasing with the house-

Figure 3.1: *Composition of household portfolios in the euro area as a percentage of total assets, by net wealth group.*



Source: Distributional Wealth Accounts.

hold's level of wealth).¹³ Risky and business assets only become prominent for households in the top decile, who have a much more diversified and liquid (as well as, of course, larger) portfolio. Following the categorisation adopted (among others) in [Slacalek et al. \(2020\)](#), we can only characterise these households (and perhaps the ninth decile) as 'non hand-to-mouth' (i.e., wealthy and with significant liquid wealth). Conversely, on average, all other groups have for the most part non-liquid wealth (housing) and limited liquid assets that are often insufficient to cover all outstanding debt. Due to the constrained nature of their balance sheets, households in the bottom 80-90% of the wealth distribution may not be able to respond to shocks as optimally as households in the top 10%. Results from our analysis supports this hypothesis.

Finally, it is striking to see how little the relative shares of different assets in households' portfolios have changed over the period covered by the DWA.¹⁴ Indeed, with the exception of a stark decrease in debt for households in the bottom 50% of the wealth distribution (partially as a result of the deleveraging process that followed the Global Financial Crisis and the euro area sovereign debt crisis), it is remarkable how the portfolio shares remained relatively stable over the years, in spite of the sometimes large volatility of some of the underlying assets (namely equity and house prices). This suggests that households' preferences in terms of saving and investment have not changed noticeably since 2009/2010. While our results show that monetary policy shocks do appear to lead to significant changes in households' portfolio composition in the short to medium term, these effects would seem to cancel each other out over longer horizons.


Figure 3.2 (see Figure E.2 for an equivalent with all asset types) displays the actual monetary values of portfolio holdings by wealth group. In addition to the considerations about portfolio compositions already pointed out above, we may also note that the overall increase in real¹⁵ wealth across groups, is largely driven by the evolution of housing. However, while the net wealth of households in the upper half of the wealth distribution increased as a result of rising asset values, the net wealth of those in the the lower half rose due to a reduction in their outstanding debt. Significant

¹³And even then, these households' exposure to risky assets is mostly due to insurance instruments rather than equity (see Figure E.1). Both insurance and mutual funds, while classified as 'risky' here, have large exposures to risk-free financial assets.

¹⁴Figure 3.1 displays results for the whole euro area. While on a country-by-country basis the portfolio shares are generally more volatile, this fact still holds in most cases.

¹⁵Figure 3.2 shows asset values adjusted for inflation. This is also the case throughout the remainder of the analysis.

Figure 3.2: *Composition of household portfolios in the euro area, by net wealth group (millions of 2015 €).*



Source: Distributional Wealth Accounts.

inequality can also be seen from the chart. Although the bottom half of the wealth distribution has five times more people, their combined net wealth is comparable to that of the sixth decile, which has lower assets but also less debt. The wealth gap widens significantly at the top: the tenth decile owns more than three times the wealth of the ninth decile, which in turn has only about 50% more wealth than the eighth decile. Overall, wealth is heavily concentrated at the top, with the bottom half owning very little in comparison.

Figure 3.3 clearly shows that over half of total net wealth in the euro area is held by the wealthiest 10% of households. The Figure also shows the extent to which assets are unequally owned across wealth groups. For instance, housing assets and non-risky financial assets are distributed much more equally across the whole wealth distribution than business assets and risky financial assets. Households in the top wealth decile own the overwhelming majority of the latter two, due to the very low participation of poorer households in these assets. While significant movements in the relative shares of holdings of certain assets can be seen from Figure 3.3 (for instance, the decrease of the overall share of debt in the bottom 50% of the wealth distribution, and the increase of the share of risky financial assets held by the wealthiest 10% of households), the relative shares of overall net wealth among household groups have remained quite steady over time.


3.3. THE EMPIRICAL APPROACH

In order to assess the impact of monetary policy on the asset levels and portfolio allocation of households along the wealth distribution, and given the nature of our data (relatively short time series, but for several different countries), we estimate impulse responses to an identified monetary policy shock by way of panel local projections. This approach better suits our purposes compared to any VAR estimation, given the shortness of the DWA time series and our interest in quantifying the response of households' wealth to monetary policy, rather than the inference of any particular coefficient.¹⁶

We make use of the euro area monetary policy shock series as constructed in Jarocinski and Karadi (2020) ('J&K'). These shocks are estimated based on asset movements during a short time window around the announcement

¹⁶See Jordá (2005), Jordá and Taylor (2024) and Plagborg-Møller and Wolf (2021).

Figure 3.3: Household group's share of total holdings in each asset type, euro area.



Source: Distributional Wealth Accounts.

of ECB monetary policy decisions (such that they can only be explained by the announcements themselves, and not other factors).¹⁷ These are further corrected for movements that can be related to central bank *information* shocks also conveyed together with the monetary policy decision. Thus, they can be taken as *pure* monetary policy surprises.¹⁸ To fit the monthly frequency of the J&K shock to the quarterly frequency of the DWA data, we create a quarterly series by summing the monthly J&K shocks across the quarter.

As our goal is to understand how the different assets and liabilities of each wealth group respond to monetary policy shocks, we estimate a set of local projections for each combination of asset type and wealth group. The DWA provides for several asset classes, however some of these show limited movement and/or are only tiny components of households' portfolios; they are however expected to behave similarly to other types of assets sharing similar characteristics. For this reason and to avoid overfitting the data, we grouped assets into five broader categories as explained in Section 3.2

The econometric specification of our local projection estimates is as follows:

$$\hat{y}_{c,g,t+h} = (y_{c,g,t+h} - y_{c,g,t-1}) = \alpha_c + \beta_h s_t + \delta_h (s_t \times ds_t) + \gamma_h \mathbf{X}_{c,t} + v_{c,g,t+h} \quad (3.1)$$

Where: (i) $\hat{y}_{c,g,t+h}$ is the dependent variable representing the change in the specific series under investigation (e.g., the the (log) holdings of a certain type of asset, or the household group's share of the country's holdings in a specific wealth instrument, etc.) at horizon h for household group g in country c ; (ii) α_c are country fixed-effects; (iii) s_t is the (quarterly) J&K euro area monetary policy shock; (iv) ds_t is a dummy variable equal to 1 if the shock is negative (i.e., expansionary); and (v) the $\mathbf{X}$ matrix represents country-specific controls including the unemployment rate, log real GDP, and the difference in yields between 10-year and 3-month Government bonds. All control variables enter the equation contemporaneously (with the exception of yields) and with four lags. We further include a time trend and a quarterly dummy for each quarter between 2020-Q1 and 2021-Q2

¹⁷A similar high-frequency approach is used by Altavilla et al. (2019) and Cesa-Bianchi et al. (2020) among plenty others, and was first introduced by Kuttner (2001) and Gürkaynak et al. (2005).

¹⁸See also Kerssenfischer (2022) for a treatment of ECB information surprises versus pure policy surprises.

to cover for COVID-19. We include four lags of both the shock and the dependent variables, and estimate local projections for a horizon of 12 quarters.¹⁹ Our panel consists of all euro area countries (except Croatia), and has unbalanced quarterly frequency (while country series start from as early as 2009-Q3 and as late as 2017-Q3, the majority starts at or before 2010-Q4).²⁰

We use the interaction term on the shock to test whether monetary policy is state-dependent, i.e. whether any effects of contractionary policy are stronger or weaker than expansionary policy. This has been found, among others, in [Alpanda et al. \(2021\)](#) with respect to the business cycle, and [Furceri et al. \(2018\)](#) with respect to the monetary stance and in the context of inequality, like for our paper (with contractionary shocks having been found to increase inequality more than expansionary shocks decrease it). Regarding state-dependent local projections, [Gonçalves et al. \(2024\)](#) warn that for these estimations to be correct in the presence of a large shock, the state needs to be exogenous. In our case, the state (contractionary vs expansionary) depends on a shock that has been constructed to be exogenous to the macroeconomic situation, alleviating this concern.

In the next Section, we present and discuss the identified effect of monetary policy shocks on the following measures for euro area households: (log) levels of each wealth type (as previously classified in this Section), portfolio composition, and overall wealth inequality. Here, we consider households grouped by their position in the wealth distribution.²¹ However, since our data also allows us to perform the same analysis by considering the employment or tenancy status of households, we also briefly discuss results for these additional breakdowns in the Appendix.

¹⁹Given the limited length of our time series, a shorter horizon might be warranted for fear of losing important degrees of freedom or significant dates (notably, the strong monetary policy shocks at the end of the sample in 2022-2023). However, estimating the same local projections with only 8 or 6 quarters of horizon does not materially change the impulse response functions or confidence intervals in the common horizon periods, while interesting movements are still captured in the third year of the IRF horizon.

²⁰DWA data is subject to ongoing updates and revisions. Our analysis is based on the most recent version released by the ECB at the time of writing.

²¹We are aware of the fact that households can, and generally do, transition from one group to another over time, and that these composition changes may affect the interpretability of our results. However, due to typically limited movements in wealth, these transitions should be relatively infrequent, especially over shorter time horizons, which are the scope of our analysis. Also, due to the wealth quantiles being of constant size relative to the population, as many households need to enter a group as those that exit it, and there needs to be balance in the direction of movements; this means that households transitioning from a group to another should be similar to each other.

The impulse response functions ('IRF') reported in the following section display changes following a one-standard-deviation contractionary shock. We have standardised the original shock series to have standard deviation of 1. We report confidence intervals computed using Driscoll-Kraay (Driscoll and Kraay, 1998; Hoechle, 2007) standard errors.

3.4. MONETARY POLICY AND WEALTH DYNAMICS ACROSS THE WEALTH DISTRIBUTION

3.4.1. Levels


In this first set of local projections, the dependent variable is the logarithm of the level of wealth held in each asset type. The impulse responses displayed in Figures 3.4, 3.5, 3.6, E.3, E.4, and E.5 show the cumulative differences in these log levels at each projection horizon (i.e., the approximate per cent growth rate) in response to the shock for households in each wealth group.

Broadly in line with the different channels of monetary policy outlined above, a number of assets display a decline for most wealth deciles (whether through price effects or changes in saving behaviour, or a combination of the two). However, these movements significantly differ in intensity across wealth types and household groups.

Non-risky financial assets (Figure 3.4) decline significantly and consistently across household groups, which is consistent with a contraction in the monetary aggregates suggested by theory. The decline is very pronounced and increases throughout the horizon, but only becomes statistically significant towards the end. The response is also strikingly similar: all groups see non-risky financial assets decline by between 2 and 3% (with some variation) except for the top 10% of the wealth distribution, where these decline by less and not significantly. While we see the channels of monetary policy working in the case of safe assets, we observe no clear response when it comes to debt (Figure E.3).

The response of risky assets (Figure E.4) is more volatile and shows more variation across the wealth distribution. While rarely statistically significant, we observe substantial declines in these assets especially for deciles 6, 7 and 8. These households and the bottom 50% have only very limited exposure to this type of assets, hence such movements do not excessively affect their overall wealth. For decile nine and the top 10%, there is an

Figure 3.4: *IRFs: cumulative percent changes in households' non-risky financial assets following a one-standard-deviation contractionary monetary policy shock, by net wealth group*




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

initial strong increase peaking at about the middle of the horizon, which then dissipates. For the top decile in particular, the increase is very large and significant two years after the shock; as they own the largest share of risky assets by far, this might appear slightly at odds with the expectation of monetary policy negatively affecting equity prices. This might be due to investment behaviour differing for this group of households.

Compared to financial assets, the response of housing (see Figure 3.5) to monetary policy is not as large in percentage points and is never significant. This is not surprising, given that housing is a very illiquid asset. However, due to the predominance of this asset within household portfolios across the whole wealth distribution, its movements are still a major driver of net wealth developments. An initial (albeit non-significant) decline observable (of around 1%) in most groups after about a year is consistent with what we separately estimate the effect of the shock to be on housing prices; this however goes back to zero soon afterwards.

Finally, we find that the response of business assets (Figure E.5) is again heterogeneous, and quite puzzling to interpret, with a very strong and positive response of decile 7 in particular, but also deciles 6, 8, 9. Not only are these responses surprisingly strong, they are also in the opposite direction relative to our prior expectations. However, due to the fact that

Figure 3.5: *IRFs: cumulative percent changes in households' housing assets following a one-standard-deviation contractionary monetary policy shock, by net wealth group*




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

these groups have very limited holdings of these types of assets (over four fifths of which are held by the top decile), and the difficulties in estimating statistics on business assets, we would not put much weight on these results.

The picture for net wealth (Figure 3.6) is a combination of the movements in assets discussed above, and the importance of these assets in each group's portfolio. Out of the noise of the group-asset responses, we can observe here that there is heterogeneity in how groups are affected by monetary policy, especially between the top and the bottom of the distribution. Because of ambiguous movements in housing, coupled with declines in financial assets, overall net wealth remains largely unaffected by monetary policy for deciles six to nine - although statistical significance aside, deciles 6 to 8 show a moderate fall. Households in the bottom 50% of the distribution, meanwhile, display an important and significant reduction in their wealth, by over 2% by the end of the horizon. Although the increase in wealth for the top decile is only marginally significant towards the middle of the horizon, it is in the opposite direction compared to all other groups, suggesting a different behavioural response of this group. There is plenty of evidence in the literature of the different behaviour of top wealth holders: for instance, [Martínez-Toledano \(2020\)](#) finds that the

Figure 3.6: *IRFs: cumulative percent changes in households' net wealth following a one-standard-deviation contractionary monetary policy shock, by net wealth group*



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

wealthiest households are better able to optimally reshuffle their portfolio during housing boom and bust cycles.

3.4.2. Portfolio composition


Our above results show that individual asset classes evolve differently in response to a monetary policy shock. What happens then to their weight in households' portfolios? To investigate this, we use the portfolio share of each wealth type²² in our local projections specification. The resulting impulse responses are displayed in Figures 3.7, E.6, E.7, E.8, and E.9.

The heterogeneous extent to which the levels of wealth held in the different asset classes respond to (contractionary) monetary policy shocks, shown in Section 3.4.1, already gives a strong indication as to what happens to households' portfolio composition. Because non-risky financial assets decline so strongly and broadly, their share of households' total assets is similarly reduced for all groups. Conversely, the value of the stock of housing assets tends to move much less, meaning that movements in housing assets' share of households' portfolios is mainly dictated by movements in other asset classes. This is evident, for instance, for the

²²Essentially what is displayed in Figure 3.1 for the euro area.

top decile, where housing's share of the portfolio moves opposite to risky financial assets.

Figure 3.7: IRFs: change in the share of non-risky financial assets in households' portfolios following a one-standard-deviation contractionary monetary policy shock, by net wealth group




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

We again do not get a clear picture with respect to debt, but leverage appears to increase marginally for households in the lower 70% of the wealth distribution. Comparing with the results in the section above, this seems to be driven by movements in other asset types rather than an actual increase in debt. Finally, the behaviour of risky financial assets and business assets differs across groups. The IRFs are in line with the level changes described in Section 3.4.1, although we note that *a priori* we would have expected risky assets to fall more substantially, given how higher interest rates on safe assets would lower the attractiveness of riskier assets. Indeed, the very sizeable increase of these assets in the portfolio of the top 10% is striking. Finally, we reiterate the caveat about business assets mentioned above.

3.4.3. Wealth inequality

We have seen above how both the levels and the portfolio shares of certain assets evolve in response to monetary policy shocks, and how these movements differ across wealth groups. What are the implications of these facts on wealth inequality? To investigate this, we use the changes in each

Figure 3.8: IRFs: household group's share of country's net wealth following a one-standard-deviation contractionary monetary policy shock




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

household group's share of the country's total holdings of a certain asset (i.e., the values displayed in Figure 3.3) as dependent variables in our local projection equation. The results of this exercise are shown in Figures 3.8, E.11, E.10, E.12 and E.13.

Throughout our analysis in the sections above, we have noted how the movements in wealth of the top wealth holders can differ substantially from the rest of the population in a number of cases. It is therefore no surprise to find that the share of total net wealth held by households in the top 10% of the wealth distribution significantly changes in response to a restrictive monetary policy shock. More specifically, as displayed in Figure 3.8, it increases. This seems to come primarily from this group's increase in the country's share of both risky and non-risky financial assets, while we see no significant movements in the distribution of housing assets across groups, apart from an increase for the ninth decile. Equally, no clear pattern is visible for debt.

Conversely, all other household groups see their share of net wealth fall. Households in the lower half of the wealth distribution are the most affected: after 3 years from the shock, their share of net wealth is about 0.2% lower; other groups' shares fall by a more moderate amount and behave similarly to each other, while at the peak after 8 quarters, the top

Figure 3.9: IRFs: household group's share of country's non-risky financial assets following a one-standard-deviation contractionary monetary policy shock



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.


10% gains around 0.5%. While the magnitudes of these movements may appear small, they are substantial, especially when compared with the very limited movements in household groups' wealth shares across the time series.²³ The strongest increase in concentration occurs for risky financial assets (which as seen are already predominantly held by the top 10% in the first place) - the top decile's share of these assets increases by almost a full percentage point. Non-risky financial assets follow a similar pattern.

As wealth becomes more concentrated, the overall level of wealth inequality therefore increases in response to contractionary monetary policy (and so does the inequality of holdings in a number of asset classes). What about expansionary shocks? Recall that in our equation, we included an interaction term between the shock and a dummy indicating whether the shock is expansionary.

While expansionary monetary policy shocks do have an opposite effect to contractionary shocks, this is not symmetric: the coefficients on the interaction term of the state dummy with the shock are negative and often significant, meaning that the reduction in inequality following expansionary shocks is less than the increase following contractionary shocks. We

²³Consider for instance that the bottom 50% own around 5% of net wealth in the euro area in our sample; to see this share decline by 0.2 percentage points is quite substantial - it is 4% of the initial share of 5%.

Figure 3.10: IRFs: household group's share of country's net wealth following a one-standard-deviation contractionary monetary policy shock (blue lines) vs a one-standard deviation expansionary monetary policy shock (red lines)



Notes: Dashed lines: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

show this in Figure 3.10, where the decrease in the top 10%'s share of wealth in response to an expansionary shock (red line) is not as large as the corresponding increase in response to a contractionary shock of the same size, and is not significant after two years. A similar story occurs with the bottom 50% of the population: their wealth share increases by less than the corresponding decrease in response to a contractionary shock. This asymmetry in the response of inequality conforms with results by Furceri et al. (2018) and points to monetary policy non-neutrality through the business cycle, as the effects of loose policy do not completely make up for the effects of tight policy.

3.5. CONCLUSIONS

In this paper, we have investigated the impact of monetary policy shocks on household wealth across the wealth distribution employing newly released, quarterly, distributional wealth data for the euro area. By using panel local projections, we have estimated the effects of monetary policy shocks on asset levels, portfolio composition, and asset concentration across the wealth distribution.

We find that contractionary monetary policy shocks lead to a decline in most asset classes for a large part of the wealth distribution, driven by

non-risky financial assets; housing assets appear affected by monetary policy to a significantly lesser extent. Households in the bottom 50% of the wealth distribution see sharper declines in their assets; at the other end of the distribution, conversely, the top 10% of wealth holders their share of asset holdings increase in response to the same shocks. As a result of these dynamics, we note an increase in wealth inequality as a result of restrictive monetary policy. We also note how the opposite occurs with expansionary policy, however the responses appear to be not as strong.

Our results are broadly aligned with previous evidence from the literature, as well as with economic theory, and complement these earlier findings by providing evidence from quarterly time-series distributional data for the euro area, for the first time.

Bibliography

- Acharya, V. V., K. Bergant, M. Crosignani, T. Eisert, and F. J. McCann (2020). The anatomy of the transmission of macroprudential policies. Technical report, National Bureau of Economic Research.
- Adam, K. and P. Tzamourani (2016). Distributional consequences of asset price inflation in the euro area. *European Economic Review* 89, 172–92.
- Aksoy, Y., H. S. Basso, R. P. Smith, and T. Grasl (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics* 11(1), 193–222.
- Alpanda, S., E. Granziera, and S. Zubairy (2021). State dependence of monetary policy across business, credit and interest rate cycles. *European Economic Review* 140(103936).
- Alpanda, S. and S. Zubairy (2016). Housing and tax policy. *Journal of Money, Credit and Banking* 48(2-3), 485–512.
- Alpanda, S. and S. Zubairy (2017). Addressing household indebtedness: Monetary, fiscal or macroprudential policy? *European Economic Review* 92, 47–73.
- Altavilla, C., L. Brugnolini, R. S. Gürkaynak, R. Motto, and G. Ragusa (2019). Measuring euro area monetary policy. *Journal of Monetary Economics* 108, 162–179.
- Alvaredo, F., A. B. Atkinson, T. Piketty, and E. Saez (2024). World Inequality Database. <http://wid.world/data>. *WID.world*.
- Amberg, N., T. Jansson, M. Klein, and A. R. Picco (2022). Five facts about the distributional income effects of monetary policy shocks. *American Economic Review: Insights* 4(3), 289–304.
- Ampudia, M., D. Georgarakos, J. Slacalek, O. Tristani, P. Vermeulen, and

- G. L. Violante (2018). Monetary policy and household inequality. *European Central Bank Working Paper* (2170).
- Antunes, A. and V. Ercolani (2020). Intergenerational wealth inequality: The role of demographics. *Banco de Portugal Working Papers* 2020(9).
- Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review* 109(6), 2333–67.
- Bartels, C. and C. Schröder (2020). The role of rental income, real estate and rents for inequality in germany. Technical report, Working Papers.
- Bernanke, B. S. and M. Gertler (1995). Inside the black box: The credit channel of monetary policy transmission. *The Journal of Economic Perspectives* 9(4), 27–48.
- Blatnik, N., A. Bobasu, K. Georgi, and M. Tujula (2024). Introducing the Distributional Wealth Accounts for euro area households. *European Central Bank - Economic Bulletin* 2024/5.
- Börsch-Supan, A., F. Heiss, A. Ludwig, and J. Winter (2003). Pension reform, capital markets and the rate of return. *German Economic Review* 4(2), 151–181.
- Carozzi, F., C. A. Hilber, and X. Yu (2024). On the economic impacts of mortgage credit expansion policies: evidence from help to buy. *Journal of Urban Economics* 139, 103611.
- Carroll, C., J. Slacalek, K. Tokuoka, and M. N. White (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics* 8, 977–1020.
- Carvalho, C., A. Ferrero, and F. Nechio (2016). Demographics and real interest rates: Inspecting the mechanism. *European Economic Review* 88, 208–226.
- Casiraghi, M., E. Gaiotti, L. Rodan, and A. Secchi (2018). A "reverse Robin Hood"? The distributional implications of non-standard monetary policy for Italian households. *Journal of International Money and Finance* 85, 215–35.
- Castellanos, J., G. Paz-Pardo, and A. Hannon (2024). The aggregate and distributional implications of credit shocks on housing and rental markets. Available at SSRN 4889614.

- Causa, O., N. Woloszko, and D. Leite (2019). Housing, wealth accumulation and wealth distribution: Evidence and stylized facts.
- Cesa-Bianchi, A., G. Thwaites, and A. Viccondoa (2020). Monetary policy transmission in the United Kingdom: A high frequency identification approach. *European Economic Review* 123(103375).
- Christophers, B. (2021). A tale of two inequalities: Housing-wealth inequality and tenure inequality. *Environment and planning A: economy and space* 53(3), 573–594.
- Cloyne, J. and P. Hürtgen (2016). The macroeconomic effects of monetary policy: A new measure for the United Kingdom. *American Economic Journal: Macroeconomics* 8(4), 75–102.
- Colciago, A., A. Samarina, and J. de Haan (2019). Central bank policies and income and wealth inequality: a survey. *Journal of Economic Surveys* 33(4), 1199–1231.
- Cooley, T. F. and E. Henriksen (2018). The demographic deficit. *Journal of Monetary Economics* 93, 45–62.
- Cooper, I. A., P. M. Sercu, and R. Vanpee (2013). The equity home bias puzzle: A survey. *Foundations and Trends in Finance* 7(4).
- David, J. M., E. Henriksen, and I. Simonovska (2014). The risky capital of emerging markets. *NBER Working Paper* (20769).
- Davtyan, K. (2017). The distributive effect of monetary policy: The top one percent makes the difference. *Economic modelling* 65, 106–118.
- De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies* 71(3), 743–768.
- Doepke, M. and M. Schneider (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy* 114(6), 1069–97.
- Doepke, M., M. Schneider, and V. Selezneva. Distributional effects of monetary policy.
- Drechsler, I., A. Savov, and P. Schnabl (2017). The deposits channel of monetary policy. *The Quarterly Journal of Economics* 132(4), 1819–1876.
- Driscoll, J. C. and A. C. Kraay (1998). Consistent covariance matrix estimation with spatially dependent data. *Review of Economics and Statistics* 80(4), 549–560.

- Dustmann, C., B. Fitzenberger, and M. Zimmermann (2022). Housing expenditure and income inequality. *The Economic Journal* 132(645), 1709–1736.
- Ehrmann, M. and M. Fratzscher (2004). Taking stock: monetary policy transmission to equity markets. *Journal of Money, Credit and Banking* 36(4), 719–737.
- Eichengreen, B. (2015). Secular stagnation: The long view. *American economic Review* 105(5), 66–70.
- Fair, R. C. and K. M. Dominguez (1991). Effects of the changing u.s. age distribution on macroeconomic equations. *The American Economic Review*, 1276–1294.
- Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The next generation of the penn world table. *American Economic Review* 105(10), 3150–3182.
- Feiveson, L. and J. Sabelhaus (2019). Lifecycle patterns of saving and wealth accumulation. *FEDS Working Paper* 2019(10).
- Ferreira, C., J. M. Leiva, G. Nuño, a. Ortiz, T. Rodrigo, and S. Vazquez (2024). The heterogeneous impact of inflation of households' balance sheets. *Banco de España Documentos de Trabajo* (2403).
- Ferrero, G., M. Gross, and S. Neri (2019). On secular stagnation and low interest rates: Demography matters. *International Finance* 22(3), 262–278.
- Franconi, A. and G. Rella (2023). Monetary policy across the wealth distribution. *Stone Center on Socio-Economic Inequality Working Paper Series* (70).
- Furceri, D., P. Loungani, and A. Zdzienicka (2018). The effects of monetary policy shocks on inequality. *Journal of International Money and Finance* 85, 168–186.
- Georgescu, O. M. and D. V. Martín (2021). Do macroprudential measures increase inequality? evidence from the euro area household survey.
- Geppert, C. (2015). On the distributional implications of demographic change. *Beiträge zur Jahrestagung des Vereins für Socialpolitik 2015: Ökonomische Entwicklung - Theorie und Politik - Session: Labor 4*.
- Geppert, C., A. Ludwig, and R. Abiry (2016). Secular stagnation? growth, asset returns and welfare in the next decades: First results. *Max Planck Institute for Social Law and Social Policy Discussion Paper* 2016(5).

- Gonçalves, S., A. M. Herrera, L. Kilian, and E. Pesavento (2024). State-dependent local projections. *Journal of Econometrics* 244(105702).
- Goodhart, C. A. and M. Pradhan (2017). Demographics will reverse three multi-decade global trends. *BIS Working Paper* (656).
- Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica* 70(1), 47–89.
- Greenwald, D. L. and A. Guren (2021). Do credit conditions move house prices? Technical report, National Bureau of Economic Research.
- Greenwald, D. L., M. Leombroni, H. Lustig, and S. Van Nieuwerburgh (2021). Financial and total wealth inequality with declining interest rates. *NBER Working Paper* (28613).
- Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking* 1(1), 55–93.
- Hoechle, D. (2007). Robust standard errors for panel regressions with cross-sectional dependence. *The Stata Journal* 7(3), 281–312.
- Hohberger, S., R. Priftis, and L. Vogel (2020). The distributional effects of conventional monetary policy and quantitative easing: Evidence from an estimated dsge model. *Journal of Banking & Finance* 113, 105483.
- Hood, A. and R. Joyce (2013). The economic circumstances of cohorts born between the 1940s and the 1970s. *Institute for Fiscal Studies Report* (R89).
- Household Finance and Consumption Network (2022). Household Finance and Consumption Survey: Methodological report for the 2021 wave. *ECB Statistic Paper* (45).
- Human Mortality Database. Human Mortality Database. *Max Planck Institute for Demographic Research (Germany), University of California, Berkeley (USA), and French Institute for Demographic Studies (France)*.
- Ikeda, D. and M. Saito (2014). The effects of demographic changes on the real interest rate in japan. *Japan and the World Economy* 32, 37–48.
- Ioannidis, C. and A. Kontonikas (2008). The impact of monetary policy on stock prices. *Journal of Policy Modelling* 30, 33–53.
- Jarocinski, M. and P. Karadi (2020). Deconstructing monetary policy surprises - The role of information shocks. *American Economic Journal: Macroeconomics* 12(2), 1–43.

- Jordá, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review* 95(1), 161–82.
- Jordá, O., M. Schularick, and A. M. Taylor (2015). Betting the house. *Journal of International Economics* 96(Supplement 1), S2–S18.
- Jordá, O. and A. M. Taylor (2024). Local projections. *NBER Working Paper* (32822).
- Kaplan, G., K. Mitman, and G. L. Violante (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy* 128(9), 3285–3345.
- Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to HANK. *American Economic Review* 108(3), 697–743.
- Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the household level. *Journal of Monetary Economics* 91, 19–38.
- Kavonius, I. K. and U. Topo (2024). Eurooppalaista tietoa kotitalouksien varallisuuden jakautumisesta. *Kansantaloudellinen aikakauskirja* 120(2).
- Kerssenfischer, M. (2022). Information effects of euro area monetary policy. *Economic Letters* 216(110570).
- Kindermann, F. and S. Kohls (2018). Rental markets and wealth inequality in the euro-area. *University of Regensburg, unpublished*. https://economics.unihohenheim.de/fileadmin/einrichtungen/economics/Lecture_Series_Abstracts/Kindermann_rental_markets.pdf.
- Kingham, C., Y. McCarthy, and C. O’Toole (2019). How do macroprudential loan-to-value restrictions impact first time home buyers? a quasi-experimental approach. *Journal of Banking & Finance*, 105678.
- Koeniger, W., B. Lennartz, and M.-A. Ramelet (2022). On the transmission of monetary policy to the housing market. *European Economic Review* 145(104107).
- Kopecky, J. (2022). Okay boomer... excess money growth, inflation, and population aging. *Macroeconomic Dynamics* 27(6).
- Kopecky, J. and A. M. Taylor (2020). The murder-suicide of the rentier: Population aging and the risk premium. *NBER Working Paper* (26943).
- Krueger, D. and A. Ludwig (2007). On the consequences of demographic

- change for rates of returns to capital, and the distribution of wealth and welfare. *Journal of Monetary Economics* 54(1), 49–87.
- Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics* 47, 523–544.
- La Cava, G. and C. He (2021). The distributional effects of monetary policy: Evidence from local housing markets in australia. *Australian Economic Review* 54(3), 387–397.
- Lenza, M. and J. Slacalek (2024). How does monetary policy affect income and wealth inequality? Evidence from quantitative easing in the euro area. *Journal of Applied Econometrics* 2024(39), 746–765.
- Lisack, N., R. Sajedi, and G. Thwaites (2017). Demographic trends and the real interest rate. *Bank of England Working Paper* (701).
- Luetticke, R. (2021). Transmission of monetary policy with heterogeneity in household portfolios. *American Economic Journal: Macroeconomics* 13(2), 1–25.
- Maio, P. (2014). Another look at the stock return response to monetary policy actions. *Review of Finance* 2014(18), 321–371.
- Martínez-Toledano, C. (2020). House price cycles, wealth inequality and portfolio resshuffling. *WID.world Working Paper* 2020(2).
- Mavropoulos, A. and Q. Xiong (2019). Housing consumption and macro-prudential policies in europe: An ex ante evaluation. Technical report, IWH Discussion Papers.
- Mesnard, A. (2024). De nouvelles statistiques sur la distribution du patrimoine des ménages dans la comptabilité nationale. *Bulletin de la Banque de France* (250).
- Mian, A. M., L. Straub, and A. Sufi (2021). The saving glut of the rich. *NBER Working Paper* (26941).
- Mishkin, F. S. (2007). Housing and the monetary transmission mechanism. *NBER Working Paper Series* (13518).
- Mishra, A. V. (2015). Measures of equity home bias puzzle. *Journal of Empirical Finance* 34, 293–312.
- Modigliani, F. (1988). The role of intergenerational transfers and life cycle

- saving in the accumulation of wealth. *Journal of Economic Perspectives* 2(2), 15–40.
- Moreno, M. (2024). Inequality and wealth distribution among Irish households: introducing new Distributional Wealth Accounts. *Central Bank of Ireland - Quarterly Bulletin* 2024/1.
- Neri, A., M. Spuri, and F. Vercelli (2024). Distributional Wealth Accounts: methods and preliminary evidence. *Bank of Italy - Occasional Papers Series* (836).
- Papetti, A. (2021). Demographics and the natural real interest rate: historical and projected paths for the euro area. *Journal of Economic Dynamic and Control* 132(104209).
- Paul, P. (2020). The time-varying effect of monetary policy on asset prices. *The Review of Economics and Statistics* 104(2), 690–704.
- Plagborg-Møller, M. and C. K. Wolf (2021). Local projections and VARs estimate the same impulse responses. *Econometrica* 89(2), 955–980.
- Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivations and implications. *American Economic Review* 94(4), 1055–1084.
- Rubio, M. and J. A. Carrasco-Gallego (2014). Macroprudential and monetary policies: Implications for financial stability and welfare. *Journal of Banking & Finance* 49, 326–336.
- Saez, E. and G. Zucman (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics* 131(2), 519–578.
- Sajedi, R. and G. Thwaites (2016). Why are real interest rates so low? the role of the relative price of investment goods. *IMF Economic Review* 64(4), 635–659.
- Slacalek, J., O. Tristani, and G. L. Violante (2020). Household balance sheet channels of monetary policy: a back of the envelope calculation for the euro area. *Journal of Economic Dynamics & Control* 115(103879).
- Summers, L. H. (2014). U.s. economic prospects: Secular stagnation, hysteresis, and the zero lower bound. *Business Economics* 49(2), 65–73.
- Tarne, R., D. Bezemer, and T. Theobald (2022). The effect of borrower-specific loan-to-value policies on household debt, wealth inequality and

- consumption volatility: An agent-based analysis. *Journal of Economic Dynamics and Control* 144, 104526.
- Tauchen, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters* 20(2), 177–181.
- Tzamourani, P. (2021). The interest rate exposure of euro area households. *European Economic Review* 132(103643).
- United Nations Department of Economic and Social Affairs, Population Division (2019). World population prospects. (Online Edition. Rev. 1).
- Vandenbroucke, G. (2016). Aging and wealth inequality in a neoclassical growth model. *Federal Reserve of St. Louis Review Q1 2016*(98(1)), 61–80.
- Williams, J. C. (2015). Measuring the effects of monetary policy on house prices and the economy. *BIS Papers* (88).

A Appendices

A. MEASURES OF PRODUCTIVITY

The measure of total factor productivity used in the PWT dataset implies the assumption of a certain shape of the production function (see [Feenstra et al. \(2015\)](#)). Though this cannot be avoided, it could potentially bias the results of any regression including productivity, which however cannot be omitted from the analysis presented in this paper for obvious economic theory reasons. In Table [A.1](#), I repeat the analysis for specification (2) of Table [1.1](#) (as this is the specification that is later used for the analysis in the other Sections) using different indicative measures of productivity that I can construct from the rest of the data. These are:

- Output per hour worked Y/H ;
- Output per person employed Y/L ;
- A measure of joint labour and capital productivity computed as $\frac{Y}{\sqrt{K \times H}}$, which I call $Prod_3$ here, and
- A similar measure, computed as $\frac{Y}{\sqrt{K \times L}}$, which I call $Prod_4$.

I also include TFP itself and a specification without any productivity measure. The results show several things: (i) the signs of the coefficients are not affected by the choice of productivity measure; (ii) their size, however, can change significantly; (iii) expressing productivity in terms of hours worked or per person employed does not make a difference; (iv) not having any measure of productivity leads to clear omitted variable bias. The wide range of coefficients, especially for K/N , may lead to think that, potentially, the results in the rest of the paper would not be robust to a different choice of productivity measure. This is however not the case. Replacing TFP with each of these measures, I have repeated the instrumental-variable exercise of section [1.4.2](#) and derived the combined age-group implied coefficient on $\ln(r^k)$ as I have done for Figure [1.5](#) (tables

omitted for brevity). The results are displayed in Figure A.1 and show that the strong demographic effect on rates of return remains and is very similar regardless of the measure used. It also further supports the case for using the PWT TFP measure as the results using this measure are the closest to the average. If anything, based on the R^2 in specification (3) of Table A.1, $Prod_3$ might be more appropriate; because of the larger size of the coefficients, this would lead to even stronger demographic effects on the rate of return and an even larger impact of the demographic transition on wealth accumulation as shown in Section 1.6.

Figure A.1: Combined effect of each age group on rates of return via $\ln(K/N)$ and $\ln(L/N)$, by measure of productivity used in instrumental variable regression (confront Tables 1.2 and 1.3).




Table A.1: Specification (2) of Table 1.1, for different measures of productivity.

	(1) $\ln(r^k)$	(2) $\ln(r^k)$	(3) $\ln(r^k)$	(4) $\ln(r^k)$	(5) $\ln(r^k)$	(6) $\ln(r^k)$
$\ln(\alpha)$	1.012*** (0.10)	1.012*** (0.10)	1.007*** (0.05)	1.007*** (0.05)	0.930*** (0.09)	0.998*** (0.15)
$\ln(K/N)$	-0.615*** (0.09)	-0.615*** (0.09)	-0.459*** (0.06)	-0.459*** (0.06)	-0.322*** (0.05)	-0.129*** (0.05)
$\ln(H/L)$	0.918*** (0.27)	0.294 (0.24)	0.803*** (0.16)	0.250 (0.17)	0.902*** (0.26)	0.232 (0.30)
$\ln(L/N)$	0.643*** (0.14)	0.643*** (0.14)	0.442*** (0.07)	0.442*** (0.07)	0.479*** (0.13)	0.240 (0.18)
$\ln(Y/H)$	0.624*** (0.11)					
$\ln(Y/L)$		0.624*** (0.11)				
$\ln(\text{Prod}_3)$			1.105*** (0.06)			
$\ln(\text{Prod}_4)$				1.105*** (0.06)		
$\ln(\text{TFP})$					0.614*** (0.10)	
Constant	-2.345 (2.12)	-10.972*** (2.87)	-1.809 (1.70)	-9.441*** (1.49)	-3.878* (2.09)	-1.070 (2.42)
Year f.e.	✓	✓	✓	✓	✓	✓
R^2	0.651	0.651	0.880	0.880	0.656	0.541
deg. fr.	63	63	63	63	63	63
N	3000	3000	3000	3000	3000	3000

B. COMPARISON OF THE VARIOUS SIMULATIONS

Throughout the paper, we have discussed the outcomes of different variants and simulations of the model, in terms of the distributional consequences of each policy considered. In Table B.1, we summarise the most relevant distributional statistics for each of the simulations discussed above, for a quick comparison among them. We are interested in the outcomes of policies in terms of inequality (measured by Gini coefficients) in wealth, housing wealth, income and consumption, as well other as overall welfare considerations, such as the distribution of utility.

We note that the variation of housing wealth inequality across simulations is the largest, while that of overall wealth is the lowest. This suggests that increasing housing inequality is often largely compensated by lower inequality in financial wealth: the households who do not have housing save more in financial wealth instead. Income inequality (where income here is net income, i.e. after taxes and after rental income or rental payments) also has a large variation, comparable to that of housing wealth; this is in line with the fact that the variation in income across simulations is given by the different numbers and compositions of renters and landlords - indeed, housing policies have no bearing on labour income.

Speaking of rental income - another thing to note is that while (net) rental yields are what matter most for landlord households (and institutional investors) when making investment decisions, this is different for renters. Renters could be better off in an economy with a high r_h but low p_h compared to the opposite situation, if that means that the actual rental payment they make is a lower proportion of their wage in the former case. For instance, in the simulation with higher construction sector productivity and tight LTI limits, r_h is higher than under the baseline scenario (3.25% vs 3.15%); however, renters are still better off, because thanks to the significantly lower house prices (5.20 vs 5.80), their rent payments are actually lower relative to aggregate income (15.6% vs 18.3%).

In terms of utility, the largest movements across different simulations come from the poorest households, who have typically low consumption and no owned housing, and thus have very low levels of utility. Welfare considerations based on utility are tricky due to the merely ordinal nature of this measure; however, due to higher marginal utility for poorer households, 'transferring' a unit of consumption from a high-wealth, high-consumption household to a low-consumption household would increase overall welfare

Table B.1: Model outcomes, various model specifications

Simulation	Baseline	2.4.2	2.4.3	2.4.3	2.4.4	2.4.2	2.4.5	2.4.5	2.4.5	2.4.6
LTI	6	3.5	6	3.5	6	6	6	6	3.5	3.5
LTV	90%	90%	90%	90%	90%	75%	90%	90%	90%	90%
τ_h	40%	40%	40%	40%	20%	40%	40%	40%	40%	45%
Inst. inv.?	NO	NO	YES	YES	NO	NO	NO	NO	NO	YES
r^{ii}	-	-	1/3%	1/3%	-	-	-	-	-	1%
L	L	L	L	L	L	L	L*1.2	L	L	L*0.8
A_h	1.3	1.3	1.3	1.3	1.3	1.3	1.3	1.4	1.4	1.3
Housing Gini	45.2%	49.2%	44.3%	48.9%	46.5%	46.1%	46.9%	47.3%	50.7%	51.1%
Wealth Gini	44.1%	44.3%	44.8%	44.8%	44.8%	44.3%	44.9%	45.0%	45.1%	46.0%
Income Gini	35.1%	36.6%	32.0%	32.5%	34.0%	35.1%	34.3%	34.1%	35.9%	39.8%
Cons. Gini	38.3%	39.2%	35.4%	35.2%	37.0%	38.3%	36.8%	36.8%	37.6%	41.1%
Median utility	-1.134	-1.136	-1.117	-1.131	-1.134	-1.132	-1.113	-1.118	-1.125	-1.097
Avg. utility	-1.371	-1.417	-1.275	-1.278	-1.324	-1.371	-1.326	-1.329	-1.357	-1.467
Homeownership	70.0%	66.2%	67.0%	61.9%	69.1%	69.4%	68.5%	68.0%	64.6%	62.9%
p_h	5.80	5.80	5.80	5.80	5.80	5.80	5.30	5.20	5.20	6.55
r_h	3.15%	3.40%	2.40%	2.40%	2.60%	3.15%	2.95%	3.00%	3.25%	3.82%
r_h % avg. wage	18.3%	19.7%	13.9%	13.9%	15.1%	18.3%	15.6%	15.6%	16.9%	25.0%

(when seen as the sum of individual utilities). It is therefore consumption inequality, together with the homeownership rate, what matters most for the distribution of utility and welfare.¹

Homeownership rates move across simulations around the 60-70% range, with the highest actually being obtained in the baseline scenario. Because however the movements in this measure are all the result of households' optimisation,² we do not necessarily see homeownership as an inherently positive measure that should be pursued by policy, unless movements in homeownership also come together with movements in inequality and/or households' welfare. While owning a house improves utility (all else equal) from an individual household's perspective, a lower homeownership rate in the economy is in fact not necessarily bad in terms of welfare and equality. For example, under the scenario with a higher rate of land permits, the homeownership rate declines to 68.5% from the baseline's 70%; nevertheless, because of the lower house prices and rental yields, renting is much more affordable: this leads to reductions in income and consumption inequality, and better utility outcomes compared to baseline.

¹Recall that because households only gain utility from their first house, what matters is the proportion of households with at least one house, rather than the number of houses each of them owns.

²Contrast this instead with a situation in which, for instance, institutional investors are allowed to bid against households, as we briefly discussed in Section 2.4.3.

C. CHANGES IN INTEREST RATES

Due to computational challenges, interest rates in our model are set exogenously and do not originate from household maximisation and a need to balance savings and investment. Indeed, there is no capital in the model economy. Interest rates are also not a policy tool that is or should be used to influence the housing market; and finally as our model better describes the longer term, the interest rates within it can be characterised as ‘natural’ rates rather than arising from monetary policy.

However, as they still do play a very important role in the results we derive from our model, we describe here how some results may change based on their level. Indeed, the level of interest rates (note that there is a fixed wedge between the rate on deposits and the rate on loans) determines how attractive the return on housing is relative to the return on deposits.

As second and further houses do not provide additional utility to the owners, the rental yield (minus property tax, rental income tax and depreciation costs) needs to be at least as high as the return on deposits (net of capital gain taxes) for it to make it an attractive investment for landlords who have the necessary cash. The prospective landlords who would need to borrow in order to purchase additional housing would demand an even higher rate of return.³ As a result, we can expect equilibrium housing demand to decline as deposit interest rates rise while (given any level of p_h determined by the construction sector) the equilibrium rental yield would rise with interest rates.

Assume now that r_b rises to 1% from the baseline 0%. After capital gain taxes of 34%, this corresponds to a net return of 0.66%. With δ_h and τ_h totalling 1.1% between them, and a rental income tax rate of 40%, housing provides a better investment when $(1 - \tau_{rh})r_h - \delta_h - \tau_h > (1 - \tau_k)r_b$, i.e. if $r_h > 2.93\%$. This appears evident in Figure C.1, with housing demand jumping significantly as r_h surpasses 2.93%. Interestingly however, the equilibrium rental yield does not rise by as much as the increase in the deposit rate, increasing to 3.80% from 3.15% in the baseline model. Lowering the interest rate on deposits would likewise result in a lower equilibrium rental return.


The effects on inequality are similar to those seen in a number of experi-

³Some households will still demand owner-occupied housing even if deposits are more attractive financially, due to the utility gain from owner-occupied housing.

ments above: because the equilibrium r_h increases, renters end up paying higher rents, leading to higher income and consumption Gini coefficients. Landlords, however, do not really enjoy the benefits of a higher rental yield, because the mortgage interest rate r_l actually increases by more than r_h (i.e., only landlords who buy investment properties out of pocket stand to gain).

Note finally how a change in the property tax rate would lead to very similar results to a change in interest rates as described here. Indeed, a higher property tax rate decreases the attractiveness of housing as an asset by reducing one-for-one the (net) differential in returns between housing and deposits. It therefore leads to a decrease in demand for housing at any price and to an eventual increase in the equilibrium rental yield.

Figure C.1: *Surface area representation of demand for housing as a percentage of population, when r_b is increased to 1%, all else equal, by p_h and r_h .*



D. RESULTS OVER TENANCY AND WORKING STATUS BREAKDOWNS

We have shown how the impact of monetary policy on the wealth of different households depends on the size and composition of their balance sheet at the time of the shock. DWA data shows that households differ in these aspects depending on where they are in the wealth distribution. Similarly, households with different employment and home-ownership status also have heterogeneous balance sheet compositions. Thus, it is possible to apply the same empirical strategy introduced in Section 3.3 and discuss analogous results for work and tenancy breakdowns.


It should be noted that employment and tenancy breakdowns are fundamentally different than net wealth deciles. The number of households within a specific group is, in fact, not fixed, whereas net wealth deciles in each country all contain one tenth of the total. Furthermore, individual households may ‘switch’ from one group to the other and cause issues in correctly isolating the evolution of each breakdown. For instance, if a tenant household takes up debt to acquire their main residence it would be recorded as an *increase* in debt on the side of home-owners in the following period, as it would then become part of the other group (meaning that this behaviour would not be visible from the tenants group’s perspective, despite being the origin of this movement). In other words, it is partly unclear whether changes in one group are entirely due to households originally included in that very group or not.⁴ These movements could constitute a significant bias in the results presented here.

D.1. Homeownership breakdown

Households are classified as either ‘owners or part-owners’ of their main residence dwelling, or as ‘tenants’. In the euro area, homeowners account for the large majority of all wealth types (with percentages often surpassing 90%), and are significantly richer, both overall and on a per-capita basis. Furthermore, the wealth composition of the two groups differs notably. Tenants have a more diversified portfolio, as lack of housing assets prevents

⁴The same issue of households ‘moving’ from one group to the other is also present when considering net wealth deciles breakdowns. However, the constraint of having the same, fixed, amount of households in each group largely prevents this from being a relevant problem, as this only interests households that are at the very ‘edge’ of two groups, whereas it could potentially affect all households when considering employment or tenancy characteristics.

Figure D.1: IRFs: percent changes in households' net wealth following a one-standard-deviation contractionary monetary policy shock, by tenancy status



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.


the concentration of wealth in this item (as it is otherwise found on the balance sheet of home-owners). Housing still accounts for about 25% of tenant's total assets,⁵ while risky and non-risky financial assets altogether account for around 60%. The weight of business assets and debt liabilities (as a share of total assets) is similar between tenants and homeowners.

Figure D.1 shows the change in overall net wealth after a contractionary monetary policy shock, for the two groups. The responses are statistically non-significant for both groups, but seem to be larger for tenants. Indeed, the share of wealth held by owners might be increasing by the end of the horizon, but this is very small and well within the confidence bands (Figure D.2).

While the percentage movements in non-risky financial assets are similar across the two groups - and similar to the responses of the wealth groups above (a decline of about 2%) - the portfolio share of these assets for tenants falls much more strongly due to the higher weight in their portfolios to begin with (Figure D.3). Risky financial assets in both groups move similarly to the top 10% decile as described above, suggesting that it might be the richest of both homeowners and tenants to be driving the movements


⁵Meaning that some tenant households own housing assets they do not live in. Due to the limited wealth of tenants as a category overall, these assets end up having a rather significant weight.

Figure D.2: IRFs: household group's share of country's net wealth following a one-standard-deviation contractionary monetary policy shock, by tenancy status



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure D.3: IRFs: household portfolio share of non-risky financial assets following a one-standard-deviation contractionary monetary policy shock, by tenancy status



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

in risky assets.

D.2. Employment breakdown


We now provide some insight into the behaviour exhibited by euro area households broken down by their employment status. Our data distinguishes five categories, according to the working status of the main income earner of the households: employed, retired, self-employed, unemployed or ‘undefined and other’. Overall, employees and retirees account for a similar share of total net wealth in the euro area, followed by self-employed households. However, ownership rates of specific wealth instruments differ significantly: employees alone account for about 60% of total debt in the euro area (as retired households have largely repaid their mortgages by the time they stopped working), and the largest part of total business wealth is held by self-employed households.

The distribution of wealth by employment breakdown is less unequal than by net wealth decile or tenancy status, and the portfolio composition across clusters is more similar. It is therefore expected that the impulse response functions estimated by our model will exhibit somewhat similar movements for some of the groups. This is the case for employees and retirees, who have almost identical IRFs for net wealth, and experience only marginally negative and non-significant movements in wealth by the end of the horizon (Figure D.4). The self-employed, typically wealthier and with more business assets, actually see their wealth increase, while for the unemployed we observe no response for two years after the shock, and then a decline towards the end of the horizon. There is a large increase for the ‘undefined’ group, but we are not able to further elaborate on that.

Following the monetary policy shock, we find that the share of housing in portfolios tends to fall substantially for the unemployed, self-employed and retirees, as shown in Figure D.5, while the response for employees is non-significant. This type of heterogeneity in responses was not visible when households were simply broken down by wealth group, and gives extra insight into the different responses of households in different situations (with the caveat mentioned above regarding group compositions).⁶


⁶Consider also the issue of movements across groups, again. Technically, an increase in unemployment generated by monetary policy might lead to an increase in the assets of the ‘unemployed’ group simply due to the fact that the group is growing in size, while what we really would like to know is whether the households who were already

Figure D.4: IRFs: percent changes in households' net wealth following a one-standard-deviation contractionary monetary policy shock, by employment status



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.


Figure D.5: IRFs: households' portfolio share of housing assets following a one-standard-deviation contractionary monetary policy shock, by employment status



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

While as we mentioned the decline in wealth for employees and retirees in response to a contractionary shock is minimal, we can see from Figure D.6 that nevertheless both groups lose a substantial share of the country's total net wealth, at the expense of the self-employed and 'undefined' groups. While the unemployed see their net wealth decrease substantially in percentage terms, their wealth is such a small portion of the total to begin with that its effect is tiny relative to the country's total wealth.

Figure D.6: *IRFs: households' country share of net wealth following a one-standard-deviation contractionary monetary policy shock, by employment status*




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

unemployed before the shock are worse off.


E. ADDITIONAL CHARTS

Figure E.1: *Composition of household portfolios (detailed breakdown) in the euro area, by net wealth decile.*




Source: Distributional Wealth Accounts.

Figure E.2: *Composition of household portfolios (detailed breakdown) in the euro area, by net wealth decile.*




Source: Distributional Wealth Accounts.

Figure E.3: IRFs: per cent changes in households' debt liabilities following a one-standard-deviation contractionary monetary policy shock, by net wealth decile




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.4: IRFs: per cent changes in households' risky financial assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.5: *IRFs: per cent changes in households' business assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile*




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.6: *IRFs: households' debt over total assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile*




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.7: IRFs: households' portfolio share of housing assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.8: IRFs: households' portfolio share of risky assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.9: IRFs: households' portfolio share of business assets following a one-standard-deviation contractionary monetary policy shock, by net wealth decile




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.10: IRFs: household group's share of country's debt liabilities following a one-standard-deviation contractionary monetary policy shock




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.11: IRFs: household group's share of country's non-risky financial assets following a one-standard-deviation contractionary monetary policy shock




Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.12: IRFs: household group's share of country's risky financial assets following a one-standard-deviation contractionary monetary policy shock



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.

Figure E.13: IRFs: household group's share of country's business assets following a one-standard-deviation contractionary monetary policy shock



Notes: Dark-shaded area: one-standard-deviation robust confidence interval. Light-shaded area: 90% robust confidence interval. Solid line: local projection moving-average-smoothed IRF.