

University of Oxford
Department of Economics

Thesis submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

---

Essays on Demographics and Heterogeneity in Macroeconomics

Pierre E. Weil
Balliol College

Trinity 2019

STATEMENT OF ORIGINALITY

I certify that the intellectual content of this thesis is the product of my own work and that all assistance received in preparing this thesis and relevant sources have been acknowledged. This thesis has not been submitted for any other degree or purposes.

Pierre E. Weil
Oxford,
Trinity 2019

ACKNOWLEDGEMENTS

I am indebted to my supervisor, Guido Ascari, for his guidance and support throughout my doctoral research. I am especially grateful to Andrea Ferrero and Benjamin Moll for their advice and comments. Many thanks also to Rich Harrison, Michael McMahon, and Ricardo Reis. I am grateful to the members of the Department and of the Bank of England, as well as the many seminar participants whose discussions have greatly improved this work.

I can only begin to thank my parents and Rebecca for their support and encouragement on this journey. I am grateful to Balliol College and to the Manorites for providing a haven of *otium* and companionship over these years.

I gratefully acknowledge financial support from the Department of Economics, the David Walton Memorial Fund, and Balliol College.

BSD

ABSTRACT

This thesis consists of a collection of essays on demographics and household heterogeneity in macroeconomics. It is divided into three parts. The first builds a new framework for writing overlapping generations models with heterogeneity in household wealth and income, building on recent technical advances based on working in continuous time. The main insight is to treat household age in a symmetric way with other axes of household heterogeneity. Even in its simplest configuration, I show that this model produces empirically realistic distributions and lifecycle behavior.

The second part of this thesis examines the mechanisms of the model and consists of two chapters. The first provides a quantitative and theoretical analysis of the distributional implications of adding lifecycle to a heterogeneous agent framework. I show that intergenerational transmission slows down the dynamics of the wealth distribution. As a result, commonly proposed mechanisms for explaining the long right tail of the empirical wealth distribution cannot explain observed changes in wealth inequality. The second chapter examines the distributional implications of population aging and rising inequality, highlighting potential distributional feedback loops that operate through asset prices.

The final part analyzes the macroeconomic and distributional consequences of changes to redistributive fiscal policy, with a focus on pensions. Transfers crowd out private savings heterogeneously for households in different income, wealth, and age groups. These changes cumulate to dynamic effects on the wealth distribution. I use the framework of Part I to quantify these effects, making two main contributions. First, I measure the cross-sectional impact of pension reform on household savings and find that adjustment is concentrated among workers in lower–middle wealth groups. The equilibrium real rate stays largely unchanged, in support of previous partial equilibrium approaches to this problem. Second, I show that raising social security benefits leads wealth inequality to fall in the short run, but to grow past its original level after fifteen years – even if accompanied by a progressive tax increase. Means testing amplifies this effect. Progressive transfers to young workers have similar impact, but mainly through altering household portfolio choice.

CONTENTS

Contents **iv**

List of Figures **vi**

List of Tables **viii**

Introduction **1**

I **10**

1 OLG Heterogeneous Agent Models in Continuous Time **11**
1.1 Introduction 11
1.2 A Unified Model of Heterogeneity in Age, Income, and Wealth . . 20
1.3 Model Parameterization and Fit 39
1.4 Concluding Remarks 52
1.A Solution Method: Derivations and Proofs 54
1.B Solution Method: Numerical Strategy 65

II **79**

2 Fat Tails in Wealth from Fat Tails in Age? **80**
2.1 Introduction 80
2.2 Demographics and Idiosyncratic Capital Return Risk 83
2.3 Quantifying the Tail of the Wealth Distribution 90
2.4 Conclusion 93
2.A Proofs 94

3 Inspecting the Mechanism **99**
3.1 Introduction 99
3.2 Why the Lifecycle Matters 100
3.3 Population Aging and Rising Inequality 108
3.4 Distributional Effects of an Interest Rate Change 123

III **131**

4 Redistribution from the Cradle to the Grave **132**
4.1 Introduction 132
4.2 Model 138
4.3 Distributional Effects of Social Insurance 140
4.4 Conclusion 160
4.A Further Analysis and Robustness Results 162

Conclusion **165**

Bibliography **169**

LIST OF FIGURES

1 Wealth inequality by age group in the US 2
2 Inequality is rising and populations are aging in advanced economies 9

1.1 Profiles of wealth and consumption by age in the US 11
1.2 The age distribution in Japan: population aging 15
1.3 Concave consumption functions 29
1.4 Calibrated mortality rates and lifecycle wage 44
1.5 Lifecycle profiles by productivity group 48
1.6 Hump-shaped lifecycle profiles of consumption and income 50
1.7 Gini coefficients of wealth within each age group 51

2.1 Tail wealth inequality in the US 81
2.2 Distribution of wealth in the model and the data 91

3.1 Savings and wealth inequality in lifecycle and perpetual youth models 106
3.2 Population aging: the demographic shock 110
3.3 Population aging: Dependency ratio and budget pressure from population aging in transition 112
3.4 Population aging: Equilibrium interest rate and wealth inequality in transition 113
3.5 Population aging: The endogenous increase in retirement age required to maintain budget surplus 116
3.6 Population aging: The equilibrium path of the interest rate under different policies 117
3.7 Population aging: The path of the shares of wealth held by the top 10% and 1% for different policies 117
3.8 Population aging: Aggregate welfare over time for different policy responses 118
3.9 Population aging: Aggregate household wealth under different policy responses 118
3.10 Effect of higher income inequality on lifecycle consumption–savings and inequality 122
3.11 Effect of interest rate hike on lifecycle profiles 126
3.12 Bond holdings and wealth inequality after an interest rate hike 127

4.1 Increasing transfer spending and wealth inequality in the U.S. 132
4.2 Income tax profiles: baseline and raised for extra transfers 143
4.3 Changes in savings rate in response to social security increase 146
4.4 Long-run change in wealth inequality by age group after a social security increase 148
4.5 Long-run impact of social security increase on the marginal distribution of wealth 149
4.6 Dynamic response of savings, consumption, and wealth inequality to social security increase 150
4.7 Effect of means-testing social security transfers on savings and the wealth distribution, over the lifecycle 155
4.8 Effect of means-tested transfers to workers on lifecycle wealth accumulation 157
4.9 Effect of various policies on lifecycle wealth inequality 158
4.10 Dynamic effects of an increase in means-tested transfers to workers . 159

LIST OF TABLES

1.1 Estimated parameters of productivity process (annualized). 43
1.2 Fit of wealth distribution 46

2.1 Steady state wealth distribution statistics. 92

3.1 Effect of lifecycle: model aggregates and inequality. Aggregates in thousands of dollars. 104

4.1 Change of aggregates over pension wealth change, with and without corresponding tax increase 145
4.2 Reactions of aggregates to pension reform size (elasticities) 162

INTRODUCTION

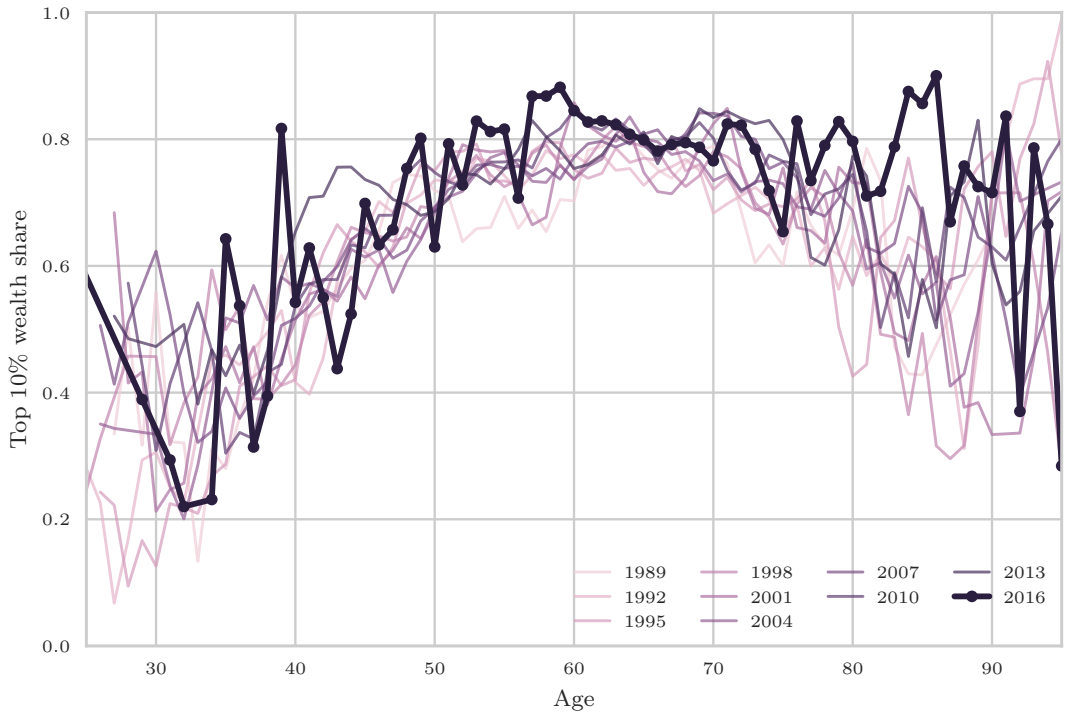
Macroeconomists have long wavered over the tradeoff between simplicity and realism in models. The intellectual history of the representative agent epitomizes this choice. These widely adopted models assume that all households are identical, replacing the wide variety of agents with an average consumer. Yet, macroeconomists have also seen the need for a sharper approach in certain settings, developing heterogeneous agent models. Two examples of this are particularly notable. The first allows for agents' decisions to vary with their age (*lifecycle* and *overlapping generations models*). In the second, households' behavior depends on their income and wealth.[1] In these contexts, it becomes crucial to keep track distributions of household characteristics such as age, wealth and income. These household attributes are very heterogeneous in the data, as illustrated by the plot of wealth inequality by age in the United States in Figure 1. Macroeconomists increasingly accept the notion that these distributions are important, even in contexts where the focus is elsewhere: Deaton (2016) writes that "except in exceptional cases, macroeconomic aggregates themselves depend on distribution."

Two global trends place renewed emphasis on household heterogeneity. These are illustrated in Figure 2. First is the *demographic transition*, with which birth rates fall and life expectancies rise as a country develops, leading to population

[1]Kaplan and Violante (2018) provide a thorough study of this literature.

Figure 1: Wealth inequality within age groups in the United States

Note: Share of age group wealth held by top 10% wealthiest in age group.
Data: Survey of Consumer Finances.



aging. Second is the striking evolution in the distributions of income and wealth of advanced economies. Inequality has risen, and earnings are increasingly polarized. For example, Piketty et al. (2017) finds "a sharp divergence in the growth experienced by the bottom 50% versus the rest of the economy": while average real income has grown by 60% since 1980, that of those in the lower half has stagnated.

These two trends share several characteristics. Changes in expected lifespan and fertility affect lifecycle patterns in household consumption, saving, and investment decisions. For example, workers need to plan for a longer retire-

ment with fewer children to support them, pushing up savings rates (Hansen 1939; Carvalho et al. 2016). Similarly, the consumption and investment behavior of households varies with their wealth and income. In particular, less well-off households tend to consume a greater share of their resources, so that rising inequality in wealth and income is associated with more aggregate saving (Fisher 1930; Keynes 1936; Dynan et al. 2004). These changes to household saving behavior affect all aspects of the modern economy. Fiscal policymakers face new pressures, both to respond to the increasing ratio of retirees to workers and to redistribute to address inequalities. Differences in risk attitudes between retirees and workers, or between rich and poor, change the relative demand for assets and affect prices. In both cases, extra savings depress the natural real interest rate, taking monetary policy closer to the zero lower bound. Both population aging and rising inequality may also be associated with the slower growth observed in advanced economies in recent years (*secular stagnation*, see Eggertsson and Mehrotra (2014) for example).

Interactions between these trends raise further economic issues. For fiscal policy, Lindert (2017) writes about the tension between these two priorities, with a tendency in advanced economies for social spending to drift towards supporting seniors at the expense of redistribution and social insurance proper. This may lead to more inequalities over the medium term. For monetary policy, the downward pressure exerted on the natural rate of interest emphasizes the risk of the zero lower bound, emphasizing consideration of unconventional policies such as quantitative easing or increasing the central bank inflation target. For welfare and inequality: wealth is a vehicle for self-insurance and for intertemporal substitution of consumption, influenced by myriad factors including income risk,

lifecycle, asset prices, and fiscal and monetary policy, each of which is affected. Interactions exist even at the household level: an extra dollar saved to finance a now longer retirement can be used in an emergency to smooth out fluctuations in earnings, modifying the demand for social insurance and the pass-through to shocks of household consumption.

Due to differing emphasis and technical constraints, macroeconomists have typically focused on one or the other form of heterogeneity, using different language and methods for modeling lifecycle behavior versus differences in income and wealth.[2] Nevertheless, as we saw, there are many similarities and interactions between these kinds of household heterogeneity. Expanding on new methods, this thesis provides a framework for solving models with households that are heterogeneous in income, wealth, and lifecycle, by treating these axes of household differences in a unified and systematic way. This approach yields computational improvements over existing methods, while retaining the ability to prove theoretical results. These models are solved numerically, building on tools developed by Achdou et al. (2017) based on modelling in continuous time. In this thesis, I pursue several applications of this framework.

Macroeconomic questions pertaining to distributions and heterogeneity can be placed into two broad classes. The first analyzes the distributional effects of specific policies or mechanisms. The other identifies the macroeconomic and policy implications of changes in distributions. This thesis mostly studies issues in the first class. In Chapter 4, I analyze the distributional impact of changes to social insurance programs. In Section 3.4, I characterize the winners and losers

[2]There are a number of notable exceptions, which I discuss in detail in the introduction to Chapter 1.

from a permanent change in the risk-free rate. However, I also provide some analyses of the second class, such as different analyses in Chapter 4 and section 3.2 of the general equilibrium implications of changes lifecycle savings motives, or an examination of the effects of rising income inequality in Section 3.3.

A broad research goal remains to develop our understanding of the macroeconomic and policy effects of distributional changes. Indeed, my objective is that this thesis provide a framework for operationalizing the study of heterogeneous agent economies with lifecycle. To that end, in forthcoming research I explain how to expand on this thesis to build dynamic stochastic general equilibrium models that feature lifecycle and heterogeneity in income and wealth. These results open the door to new workhorse macroeconomic models that can be used to address the questions raised by population aging and distributional changes.

This thesis is divided into three parts. The first consists of one chapter, and sets out the main overlapping generations heterogeneous agent framework which is applied in subsequent parts. I develop new tools for solving overlapping generations models with heterogeneity in income and wealth, both within and between different age groups. I build on methods developed by Achdou et al. (2017) for financial heterogeneity and work in continuous time. That chapter sets out the framework in a general form, and explains the solution method with the aim of allowing the macroeconomist to adapt it to different situations. This general equilibrium model features idiosyncratic risk to household labor and capital income, mortality risk, bequests, and a government implementing tax–transfer fiscal policy. In its simplest configuration, I show this model produces empirically realistic distributions and lifecycle behavior, without departing from standard permanent income consumption hypotheses.

The second part of the thesis exhibits applications of the framework to several different settings, and consists of two chapters. In Chapter 2, I examine idiosyncratic risk to capital returns, a mechanism suggested by Benhabib et al. (2011) for explaining the fat tail of the wealth distribution in the United States, and argue that it struggles to explain two important facts on the dynamics of the wealth distribution in the 20th century. Using a combination of theoretical and numerical results, I show that this mechanism is slow, operating on the order of generations. It would suggest that current levels of wealth inequality are largely explained by dynasties that are hundreds of years old – and yet that these were somehow not present in the 1970s.

Chapter 3 consists of three free-standing sections, aiming to analyze the mechanisms of the model in more detail. I quantify the contribution of lifecycle savings to aggregate household wealth, by analyzing the incremental effect of introducing lifecycle motives to a baseline model of uninsured income risk. I find that lifecycle behavior accounts for 20% of aggregate household wealth, depressing the risk-free rate by a further 150 basis points. This provides support for the idea that lifecycle behavior matters even for broad-based economic aggregates. Thus, lifecycle savings exacerbate the excess saving in models with uninsured income risk first identified by Huggett (1993) and Aiyagari (1994), in which households hold precautionary savings to insure themselves against income risk. The addition of lifecycle considerations to this setting thus speaks to asset pricing puzzles which ask why safe interest rates are so much lower than predicted by standard models. Population aging strengthens these effects, as lengthening lifespans amplify the motive to save for retirement.

I examine secular stagnation in the model, looking at the macroeconomic and

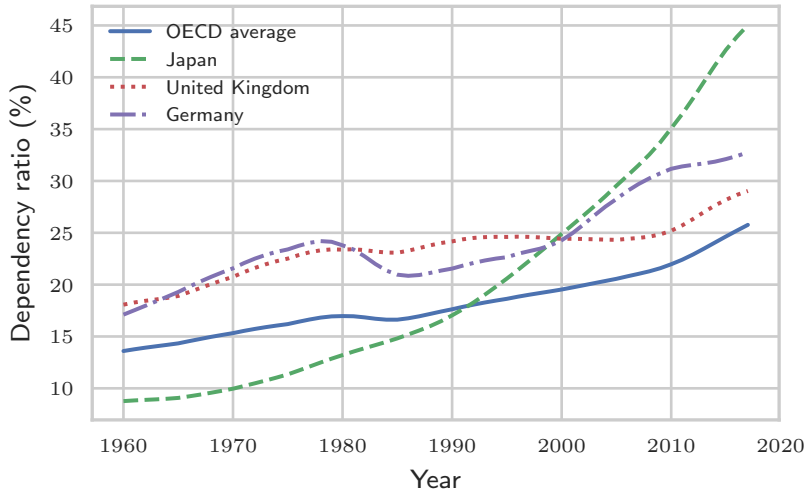
distributional implications of both population aging and rising structural income inequality. I show that these trends raise aggregate savings and depress safe interest rates in the model, and quantify the elasticities as relatively high.

Finally, I analyze the distributional effect of a permanent increase in the risk-free rate using the model. Typically, we would expect lenders to gain and borrowers to be worse off. However, in the model as in the data, agents tend to spend many of their working years as net borrowers, using leverage to buy a risky asset (housing). Lenders tend to be better off and older – however, in their youth many of them were also borrowers. Thus, over the long run, even lenders will have paid higher borrowing costs, ending up with less wealth and consumption. Wealth inequality drops in response to the rate rise, a result of portfolio rebalancing and of these heterogeneous wealth patterns in bond holdings. In rising inequalities originally caused a fall in the risk-free rate, this mechanism points to an amplification mechanism by which low rates further increase inequality.

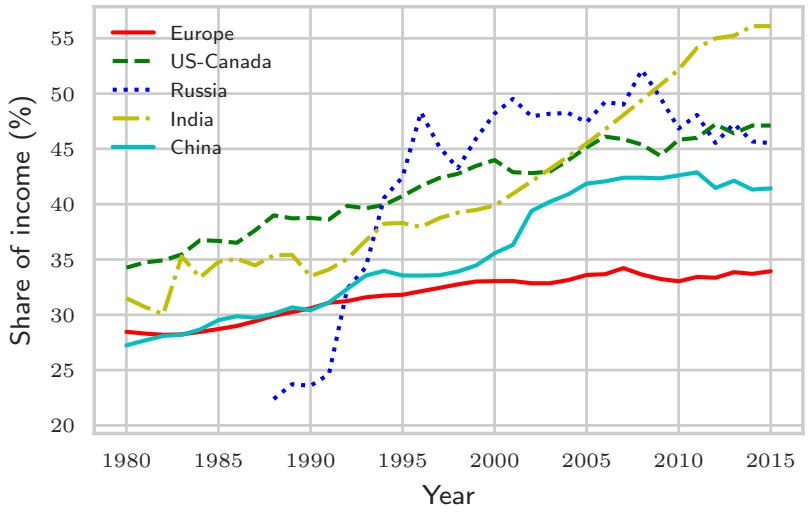
The final part of this thesis also consists of one chapter. In it, I analyze the macroeconomic and distributional consequences of changes to redistributive fiscal policy, with a focus on pensions. The evidence suggests that transfers crowd out private savings heterogeneously for households in different income, wealth, and age groups. These changes cumulate to dynamic effects on the wealth distribution, which should be taken into account for policymakers with distributional goals. In fact, more generous social insurance leads to more long-run wealth inequality, as in Hubbard et al. (1995). To quantify these channels, I use the framework of Part I and make two main contributions. First, I measure the cross-sectional impact of a pension reform on household savings. Adjustment is concentrated among workers in lower-middle wealth groups. Richer households

are indifferent about transfers, while the poorest are constrained. Thus, the equilibrium real rate stays largely unchanged, in support of previous efforts studying these effects in partial equilibrium. Second, in a transition experiment, I show that raising social security benefits leads wealth inequality to fall in the short run, but to grow past its original level after fifteen years – even if the accompanying tax increase is progressive. This is due to lower-middle class workers reducing savings most strongly. Means testing amplifies this effect. Progressive transfers to young workers have similar impact, but through different channels. Transfers encourage riskier portfolios, and crowding out is less pronounced since it is easier to save than to borrow against future transfers.

FIGURE 2: INEQUALITY IS RISING AND POPULATIONS ARE AGING IN ADVANCED ECONOMIES



(a) Ratio of elderly (65+) to working-age (18-64) population. Data: World Bank



(b) Top 10% share of income. Data: World Inequality Database 2018

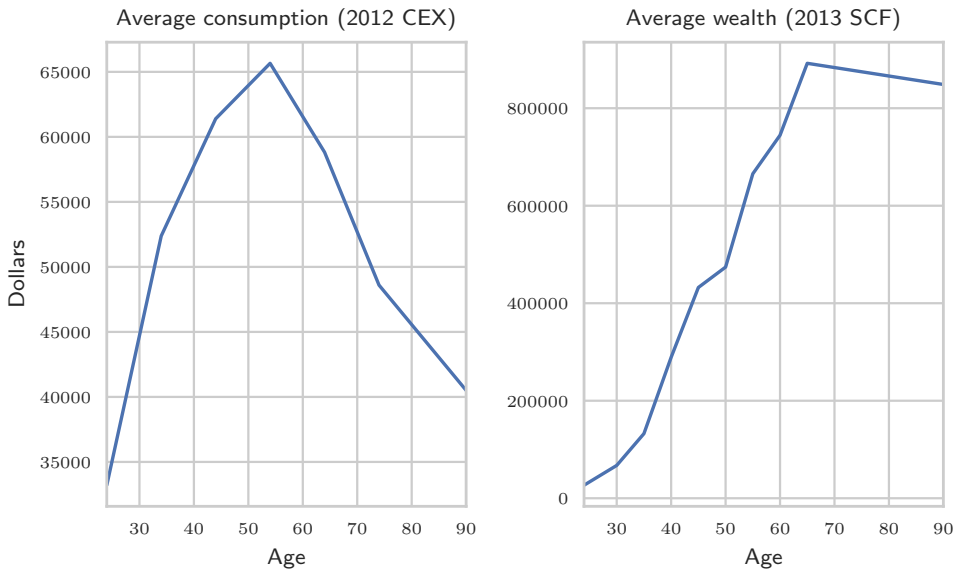
Part I

CHAPTER 1

OVERLAPPING GENERATIONS HETEROGENEOUS AGENT MODELS IN CONTINUOUS TIME

FIGURE 1.1: PROFILES OF WEALTH AND CONSUMPTION BY AGE IN THE US

Data: Consumer Expenditure Survey 2012 and Survey of Consumer Finances 2013



1.1 INTRODUCTION

How agents decide to consume and save is a central question in economics. They save for many reasons, including retirement, as a buffer against income fluctuations or unemployment (*precautionary savings*), to buy a house, to invest in a private or public business, and to leave gifts to their children. Externalities exist

between these savings motives, in the sense that once a dollar is saved it can be used for any pressing reason.

Most of the fundamental questions in macroeconomic depend on understanding the household consumption–savings decision and its implication on aggregates, asset prices, and distributions. However, for any given macro-financial environment, a household's saving decision is highly dependent on past, current, and expected future wealth and income, as well as lifecycle factors including age and family composition. As a result, addressing these macroeconomic questions requires a toolkit that can adequately model the large and increasing heterogeneity in income, wealth, and lifecycle factors that we observe in the data.

Macroeconomics has traditionally used different language and tools for studying lifecycle behavior on the one hand (*overlapping generations*) and household heterogeneity in income and wealth on the other. However, these are arbitrary distinctions. This thesis takes a systematic approach, treating both of these axes of heterogeneity in a unified way. It is my ambition that this thinking will permit new workhorse models of household heterogeneity in lifecycle and finances to join the toolkit of the macroeconomist.

In this chapter, I present a new framework for solving macroeconomic models with heterogeneity in income, wealth, and age, which I apply to several different questions later in the thesis. In particular, I develop tools for solving overlapping generations models with heterogeneity in income and wealth, both within and between different cohorts. I build on new methods for solving heterogeneous models using continuous time, expanding on tools developed by Achdou et al. (2017) for financial heterogeneity, which strike a new balance between flexibility and tractability.

The approach of this thesis is to use an overlapping generations heterogeneous agent (OLGHA) model to investigate the effects of lifecycle and distributional mechanisms in various settings. Early uses of OLGHA models include İmrohoroglu et al. (1995), Ríos-Rull (1995), and Huggett (1996), who add uninsured household income risk to large-scale overlapping generations models in the vein of Auerbach, Kotlikoff, et al. (1987), solving in steady state. More recent studies analyze equilibrium transition dynamics or aggregate fluctuations in this setting.[1] Many of these studies had to make restrictive modelling choices such as two-period lifespans or perpetual youth for technical reasons which I explain below – these can be avoided using the approach of this paper.

The technical contribution of this thesis to that literature is to provide another solution method for these types of models, using a continuous-time approach. Rather than building the large 60- or 80-period systems used previously in the literature, here the model reduces to a system of two differential equations called a *mean-field game*, which has been well studied in mathematics and physics. As a result, efficient numerical solution tools can be deployed, allowing for large state spaces of lifecycle behavior and income and wealth heterogeneity, and thus for these axes of heterogeneity to be flexibly modelled. Further, the relative tractability of this approach yields analytic results, such as the asymptotics of policy functions or the shape of the tail of the wealth distribution. This new synthesis of larger heterogeneity and improved tractability is a great advantage of the continuous time approach.

Achdou et al. (2017) derive these tools for solving models with heterogeneity

[1] These include Conesa and Krueger (1999), Castaneda et al. (2003), Krueger and Ludwig (2007), Heathcote et al. (2010), Kindermann and Krueger (2015), and Carroll et al. (2017) and the Penn Wharton Budget Model.

in income and wealth. In this paper, I extend the approach in order to solve overlapping generations models with heterogeneous households. A straightforward application of those tools runs into the *curse of dimensionality*, in which the four state variables of household income, wealth, age, and date of birth render the method computationally prohibitive. Instead, I solve the system for each cohort separately and splice these back together to find an equilibrium. This approach allows for a relatively fast numerical solution using parallelization, while retaining tractability.

In this chapter, I set out the framework in its most general setting, nesting the versions of the model used in the applications of subsequent chapters. I then calibrate a version of the model to the U.S. economy in 2013, and validate model predictions for key lifecycle and distributional moments against the data.

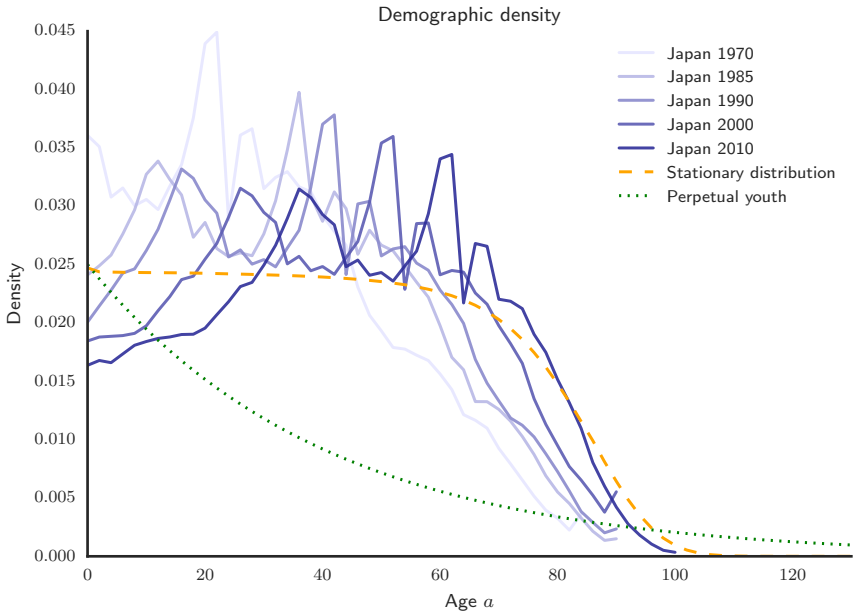
The framework, a fully continuous overlapping generations heterogeneous agent model, consists of three primary building blocks which I set out below.

1. Demographics and lifecycle: overlapping generations

Age is a crucial factor of households' economic decisions. For example, Figure 1.1 plots measures of average net worth and consumption against age in the U.S. in 2013. These are highly nonlinear, with consumption, for example, peaking around the retirement age.

To model these effects, I build on the overlapping generations literature. Early efforts to model lifecycle behavior in macroeconomics tended to rely on ad hoc approaches which permitted simple solutions. These include infinitely lived agents, finite lifespans with few periods, and *perpetual youth* – in which the probability of death is the same at every age (so called because a 20 year-old and an 80

FIGURE 1.2: THE AGE DISTRIBUTION IN JAPAN: POPULATION AGING



year-old have the same decision functions). Thus assuming away lifecycle heterogeneity in policy functions, equilibria can sometimes be solved with paper and pencil (Blanchard (1985) famously). The cost, however, is the absence of realistic hump-shaped consumption profiles over the lifecycle, as well as counterfactual demographic implications, illustrated in Figure 1.2. This figure shows the evolution of the age structure in Japan over the past decades (solid lines). The dotted line plots the age structure for a perpetual youth model calibrated for life expectancy at birth of forty years. It has a long right tail – one agent in 200 is more than 300 years old. As I discuss in Chapter 2, this can have important ramifications for the distribution of wealth in the model.

As a result, two strands of the literature grew from this point, one more computational and one more analytic. Studies such as Auerbach, Kotlikoff, et

al. (1987) solved large-period overlapping generations models numerically. By contrast, Gertler (1999) took a step towards remedying these issues by allowing for a binary lifecycle component, in which agents face two consecutive perpetual youth shocks. The first takes a working agent into retirement, and the second determines the time of death of a retiree. These methods were applied by Carvalho et al. (2016) to obtain estimates of the mechanisms and effects that demographic change has on interest rates. This family of models, consisting of two stacked perpetual youth shocks, also retain highly counterfactual demographic implications, the cost for their tractability. For example, there are a large number of workers who are centenarians.

The framework in this paper instead allows for the instantaneous probability of death to depend on age in a general way, while retaining a degree of tractability. It also allows for an arbitrary lifecycle calibration of household parameters, such as a realistic hump shape in agents' age profile of income. When agents' decision functions are different at each age, the distribution of ages becomes a key state variable. The distribution of ages, as a function from the positive real numbers to $[0, 1]$, is an (uncountably) infinite dimensional variable. By assuming away the lifecycle dimension, perpetual youth models reduce this distribution to a zero-dimensional object. In Gertler (1999), the ratio of the number of retirees to workers is a key state variable which plays the role of the age distribution; the restricted lifecycle assumptions collapse the distribution to a 1-dimensional variable.

The framework presented in this paper does not collapse the lifecycle in this way, instead allowing for arbitrarily fine specification of the lifecycle, as in Auerbach, Kotlikoff, et al. (1987) models, while allowing for a degree of tractability

present in the simpler models.

In building overlapping generations models that feature heterogeneity within age groups, the strategy has typically been to marry one of these with state-of-the-art methods in heterogeneous agent. For example, İmrohoroglu et al. (1995), Ríos-Rull (1995), Huggett (1996), and Krueger and Ludwig (2007) augment a large-scale OLG framework in the vein of Auerbach, Kotlikoff, et al. (1987) with uninsured shocks to household labor. Others, such as Castaneda et al. (2003) and Kaplan and Violante (2010), use perpetual youth. The wealth of different approaches to modelling lifecycle behavior in the literature permits macroeconomists to pick the right tool for the job. However, this can make it difficult to compare across models. By providing a framework which nests all these approaches while retaining the ability to derive analytic results, the approach of this paper can help analyze robustness of results across lifecycle modelling choices. Achieving this requires a modern solution method.

2. Solution method: heterogeneous agents in continuous time

Achdou et al. (2017) develop important new tools for solving general equilibrium macroeconomic models with heterogeneous agents. Their approach uses continuous time to reduce an economic model into a coupled system of integral and partial differential equations, forming a mathematical object called a *mean field game*. These systems capture the dynamics of a large number of small agents interacting through a distribution. Originally described in Lasry and Lions (2006, 2007), mean field games have been the subject of much research in recent years due to their wide-ranging applications to mathematics, physics, engineering and

now economics.[2] Strong existence and uniqueness results have been established for mean field games, although open questions remain. Achdou et al. (2017) provide an algorithm for the efficient numerical solution of the system. These fertile methods have led the authors to other influential papers, including Kaplan et al. (2018) and Ahn et al. (2018).

Achdou et al. (2017) is designed to analyze models of uninsured idiosyncratic risk to household income (in the vein of Bewley (1986), Huggett (1993), and Aiyagari (1994)). In this paper, I expand these methods to solve models with age-based heterogeneity (i.e. lifecycle) as well. Lifecycle effects introduce some complications in the numerical solution method which I discuss in Section 1.B. Nevertheless, solving for a stationary equilibrium takes a few seconds on a good computer, and less than an hour for a large transition.

In addition to computational speed, this continuous time approach has the advantage of retaining a degree of analytical tractability. For example, in Chapter 2 I derive theoretical results about the tail of the wealth distribution by characterizing the asymptotics of agents' policy functions. Indeed, occasionally-binding constraints are easy to handle in this setting – they bind only on the boundary of the state space. Thus, first-order conditions and usual Euler equations can easily be derived.

Wealth distribution functions are solved for directly: there is no need to simulate. Simulation can be very costly, especially when interested in modeling the tail of a distribution, since by definition these consist of very low probability events. By obtaining the whole distribution directly, this method also refines previous approaches in the literature that quantize the wealth distribution (for

[2] Achdou et al. (2014) provides a good overview of their applications to macroeconomics.

example, Benhabib et al. (2011) tracks quintiles).

As mentioned earlier, applying the tools of Achdou et al. (2017) directly in an overlapping generations setting poses a number of technical difficulties, such as increasing the dimensionality to a numerically prohibitive extent. In the appendices of this chapter I discuss in detail how to solve these issues. The algorithm and code are written in a general way, allowing for the framework to be applied to different settings without diving into the lower-level technical details.

3. Wealth and income distribution

The direct access to model distributions afforded by the solution method allows us to study the evolution of the distributions of income and wealth.

The final building block consists of the two kinds of idiosyncratic uncertainty faced by households: shocks to their labor income and shocks to their capital returns. Uninsured idiosyncratic income shocks (the *standard incomplete markets* model) lead households to build precautionary savings (Bewley 1986; Huggett 1996; Aiyagari 1994), an empirically important component of saving especially in early life (Gourinchas and Parker 2002). I will assume that household labor income has idiosyncratic fluctuations around a common age profile. These shocks will be assumed to be composed of two components, estimated by Kaplan et al. (2018) from administrative data analyzed in Guvenen et al. (2015).

Idiosyncratic risk to household capital returns, introduced by Benhabib et al. (2011), improves the modeling of the tail of the wealth distribution. The authors point out that over half of U.S. household wealth is invested in household or private business equity, both of which carry a large idiosyncratic risk component. They argue in a number of papers (Benhabib et al. 2011, 2015, 2016) that

this mechanism leads to a stationary distribution of wealth with a fat (Pareto) right tail, matching U.S. data. In Chapter 2 I qualify their findings in a life-cycle context. The authors demonstrate that idiosyncratic risk to income alone cannot account for the fat right tail in the empirical wealth distribution, as the precautionary savings motive vanishes for wealthy agents.

As a result of these shocks, the framework delivers empirically realistic distributions of income and wealth throughout the lifecycle. In later chapters, we shall use this to examine the dynamic response of these distributions to various shocks.

1.2 A unified model of heterogeneity in age, income, and wealth

This section presents a general version of the framework, nesting the formulations used in the subsequent chapters of this thesis. As these applications differ in focus, some model features will be less relevant and thus omitted.

The model runs in continuous time. A finite measure of households is formed at each point in time, with age 25. Agents have uncertain lifetimes: if alive at age t, they face a probability[3] $\lambda_t \, \mathrm{d}t$ of dying by age $t + \mathrm{d}t$. I will calibrate λ_t to match age-specific mortality rates found in actuarial tables.

The *age structure,* or density function for household age, is denoted by $D(t)$, so that $\int_0^\infty D(t) \, \mathrm{d}t$ is the total population size. I assume the entry rate is exogenous and can depend on the age structure. If (λ_t) is constant and the entry rate equals the aggregate mortality rate $\int_0^\infty \lambda_\tau D(t) \, \mathrm{d}t$, this stationary case is equivalent to the canonical perpetual youth models of Yaari (1965) and Blanchard (1985).

[3]Formally, agents die at the first arrival time of an idiosyncratic Poisson point process with rate $(\lambda_t)_{t \geq 0}$.

In an overlapping generations model, an agent's behavior depends on his age t. Decisions can vary separately with calendar time τ (or equivalently with date of birth $s = \tau - t$) if there are period (or cohort) effects. For example, $\lambda(t, \tau)$ would denote the flow death probability for agents of with age t at time τ. However, to avoid cumbersome notation, I suppress this dependence when possible.

Note that the demographic distribution evolves in an intuitive way: the agents of age t alive now are the ones who survived from a moment ago. This line of argument characterizes D by means of a partial differential equation. If $D(t, \tau)$ is the time-τ age structure, then

$$\partial_\tau D(t, \tau) = -\partial_t D(t, \tau) - \lambda(t, \tau) D(t, \tau) \qquad \text{for all } t > 0,$$

see Section 1.A.3 for details.

1.2.1 *Agent decisions*

In addition to length of life, agents face idiosyncratic risk to labor productivity and to returns on wealth. Two principal sources drive lifecycle behavior: uncertain lifespans, and an age-varying profile of average income that displays a hump shape. Labor is exogenous and takes the form of a stream of endowment income at each age.

Markets

Households choose between consumption and investment in three assets: a risk-free bond b, a risky asset k, and an annuity contract. Agents take as given the return r_t^b of the risk-free bond at each age, which is an equilibrium outcome. Following Benhabib et al. (2011, 2015, 2016), I assume the risky asset k carries

idiosyncratic risk.[4] The idea is that close to half of U.S. household wealth is tied up in housing capital or private business projects, both of which carry a large idiosyncratic component of risk. The authors argue that this mechanism is useful for explaining the fat tail of the wealth distribution observed in U.S. data – a claim which I examine in depth in Chapter 2. Assuming no portfolio adjustment, the risky asset k_t pays a return given by

$$\frac{\mathrm{d}k_t}{k_t} = r^k \,\mathrm{d}t + \sigma \,\mathrm{d}B_t. \tag{1.1}$$

The average return is r^k, with volatility σ. (Process B_t is a Brownian motion particular to each agent.) The supply of the risky asset is exogenous and infinitely elastic.

As in Blanchard (1985), agents participate in an annuity market. Each of these instantaneous contracts pays out π_t to agents aged t, in exchange for a claim of \$1 on the agent's assets were she to die at that moment. If an agent chooses to buy p_t worth of these annuities at age t, then her estate would be worth

$$q_t = w_t - p_t/\pi_t \tag{1.2}$$

were she to die, where $w_t = k_t + b_t$ is her current net worth. Agents are allowed to buy or sell annuities, with the latter case corresponding to a purchase of life insurance, but they will find it optimal to buy annuities as an insurance against mortality risk.

In a competitive market, $\pi_t = \lambda_t$. Instead, I allow for an imperfect annuity market, to match evidence on consumption patterns in old age.[5] For simplic-

[4] Aggregate risk is technically difficult to implement in these models. Ahn et al. (2018) provide promising results in this direction.

[5] Mortality rates λ_t exceed 20% for 90+ year-olds, implying counterfactually high asset

ity, the aggregate profits and losses of the annuities, as well as any accidental bequests, are consumed by the government.

Rather than p_t, I treat q_t – the total estate the agent would leave if she were to die at age t – as the agent's choice variable. Households have a bequest motive, receiving *joy of giving* (or *warm glow*) utility from their chosen bequest q_t (see, e.g. De Nardi et al. (2006) and Benhabib et al. (2016)). Agents care only about the total size of their intended legacy, not how it is distributed between multiple children nor how it will be spent.

Absent a bequest motive, households optimally choose $q_t = 0$, annuitizing all their wealth and thus never leaving accidental bequests. Otherwise, they allocate their spending between q_t and c_t. Note that $q_t \geq 0$ at all times – no insurer will sell an annuity worth more than the wealth of the agent.

Income

Agents receive exogenous income $y_t z_t$, with an age-dependent wage profile y_t common to all agents that exogenously captures lifecycle trends in labor income. This profile can be microfounded with a model of human capital accumulation. Productivity z_t carries idiosyncratic risk. Following Kaplan et al. (2018) and Guvenen et al. (2015) I assume that log productivity is mean reverting but subject to two sources of jump shocks: one small, frequent and impersistent; the other larger, rarer and more persistent. Specifically, I assume

$$\log z_t = \zeta_t^1 + \zeta_t^2; \qquad \mathrm{d}\zeta_t^i = -\zeta_t^i \,\mathrm{d}t + \mathrm{d}J_t^i, \tag{1.3}$$

returns for elderly households if $\pi_t = \lambda_t$. Annuity contracts do exist in real life; it remains puzzling that they are so rarely taken up.

where J_a^i are (independent) jump shocks, compound Poisson processes with stochastic jump sizes.

The tax and transfer schedule is captured by the function $\Theta(t, w, z)$, which can depend on household age and be asset and income tested. This general formulation captures cash transfers, income taxes, wealth taxes, and taxes to capital returns. Agents internalize the transfer schedule. Typically, Θ will include income taxes, social security, and some progressive transfers to unproductive agents. The exact specification will vary according to the different applications in later chapters.

Optimization

Agents choose lifecycle paths of consumption c_t, bequests q_t, bond holdings b_t, and risky capital holdings k_t to maximize expected discounted utility

$$\mathbb{E}\left[\int_0^\infty e^{-\rho t - \int_0^t \lambda_s \,\mathrm{d}s}\big(u(c_t) + \lambda_t \tilde{u}(q_t)\big)\,\mathrm{d}t\right], \tag{1.4}$$

where $\rho > 0$ is the agent's time discount rate. Note that agents discount consumption at any age t by the probability of not surviving to t. The utility function u for consumption has constant relative risk aversion,

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}. \tag{1.5}$$

and I assume the joy of giving utility function for bequests is given by

$$\tilde{u}(q) = \chi \frac{(q-\underline{q})^{1-\tilde{\gamma}}}{1-\tilde{\gamma}}. \tag{1.6}$$

Parameter $\chi > 0$ captures the relative strength of the bequest motive. Here $\underline{q}$ and $\tilde{\gamma}$ are variables that make the bequest motive non-homothetic when $\underline{q} \neq 0$

or $\tilde{\gamma} \neq \gamma$. In these cases, richer households have relatively stronger bequests motives, encouraging them to save more.[6]

Households maximize equation (1.4) subject to the budget constraint, expressed in terms of net financial wealth $w_t = b_t + k_t$ by substituting in equation (1.1),

$$\mathrm{d}w_t = (y_t z_t + \Theta(t, w_t, z_t) + r_t^b w_t + (r^k - r_t^b)k_t - c_t - \pi_t(q_t - w_t))\,\mathrm{d}t + \sigma k_t\,\mathrm{d}B_t, \quad (1.7)$$

and subject to the dynamics of productivity, equation (1.3).

Households face two constraints. First, they cannot hold short positions on the risky asset, so $k_t \geq 0$, a constraint motivated by informational considerations. Second, they face a borrowing constraint $b_t \geq \Phi$, where $\Phi \leq 0$. The borrowing constraint can be thought of as a limit on leverage for risky capital, leading to an upper bound on capital holdings $k_t \leq w_t - \Phi$.

The households' problem is solved recursively. Denote by $V(t, w, z)$ the optimal value of having wealth w and productivity z at age t,

$$V(t, w, z) = \max_{(c_s, k_s, q_s)_{s \geq t}} \mathbb{E}_t \int_t^\infty e^{-\rho(s-t) - \int_t^s \lambda_\xi \,\mathrm{d}\xi} \big(u(c_s) + \lambda_s \tilde{u}(q_s)\big)\,\mathrm{d}s, \quad (1.8)$$

maximizing subject to the dynamics for wealth and productivity outlined above (equations (1.3) and (1.7)), the constraints on k and b, and to productivity and net worth at age t being equal to the given values z and w.

The value function $V(t, w, z)$ satisfies the *Hamilton–Jacobi–Bellman equa-*

[6] Straub (2017) demonstrates that non-homothetic bequest motives are important for matching the concave consumption functions out of permanent income that are observed in the data.

tion:[7]

$$
\begin{aligned}
(\rho + \lambda_t)V(t,w,z) = \max_{c,k,q} \, & u(c) + \lambda_t \tilde{u}(q) + \lambda_t \tilde{u}(q) + \partial_t V \\
& + \partial_w V \big[y_t z + \Theta(t,w,z) + r_t^b w + (r^k - r_t^b)k - c + \pi_t(q-w) \big] \\
& + \partial_{ww} V \frac{\sigma^2}{2} k^2 \\
& + \partial_z V \mu^z(z) + \eta \int_{-\infty}^{\infty} (V(t,w,z') - V(t,w,z)) \phi^z(z') \, \mathrm{d}z',
\end{aligned}
\tag{HJB}
$$

where the maximum is taken over all $c \geq 0$ and $k \in [0, w - \Phi]$. See Section 1.A.1 for a heuristic derivation and a proof.

The final integral term captures the expected effect of the compound shock to productivity combining the two components, where η is the aggregate arrival rate of the jumps and ϕ^z the density function.

A concept from stochastic calculus called the *infinitesimal generator* is useful for refining our intuition about this equation. This is an operator that maps a function of a stochastic process onto the conditional expectation of that function. We will use it to rewrite the HJB equation in terms of continuation value coming from expected changes in each of the household state variables. The *generator* $\mathcal{A}_t$ of the joint process (w_t, z_t) is an operator defined on functions, mapping

$$
\mathcal{A}_t f(w,z) = \lim_{\epsilon \to 0+} \frac{E[f(w_{t+\epsilon}, z_{t+\epsilon}) \mid w_t = w, \, z_t = z] - f(w,z)}{\epsilon} \quad \text{for any } f. \tag{1.9}
$$

Ito's lemma implies that $\mathbb{E}\, df(w_t, z_t) = (\partial_t f + \mathcal{A}_t f)\, \mathrm{d}t$, so that $\mathcal{A}_t f$ captures the expected change in $f(w_t, z_t)$ that comes from dynamics of (w_t, z_t). If the policy functions are optimally chosen, we can then rewrite the HJB equation as

$$
(\rho + \lambda_t)V(t,w,z) = u(c) + \lambda_t \tilde{u}(q) + \underbrace{\partial_t V + \mathcal{A}_t V}_{\mathbb{E}\, \mathrm{d}V / \, \mathrm{d}t}. \tag{1.10}
$$

[7] For a variable x, denote by ∂_x the partial differentiation operator $\partial / \partial x$.

Further, we can calculate the infinitesimal generators $\mathcal{A}_t^w$ and $\mathcal{A}_t^z$ for processes w_t and z_t respectively, and find that for any function f,

$$\mathcal{A}_t^w f = \partial_w f \left[y_t z + \Theta(t,w,z) + r_t^b w + (r^k - r_t^b)k - c + \pi_t(q-w) \right] \tag{1.11}$$

$$+ \partial_{ww} f \frac{\sigma^2}{2} k^2, \tag{1.12}$$

$$\mathcal{A}_t^z f = \partial_z f \mu^z(z) + \eta \int_{-\infty}^{\infty} (f(t,w,z') - f(t,w,z)) \phi^z(z') \, \mathrm{d}z'. \tag{1.13}$$

Consequently, $\mathcal{A}_t = \mathcal{A}_t^w + \mathcal{A}_t^z$, and we can understand the last two lines of equation (HJB) as continuation value coming from changes in the state variables w_t and z_t respectively, and we can rewrite equation (HJB) as

$$(\rho + \lambda_t) V(t,w,z) = u(c) + \lambda_t \tilde{u}(q) + \partial_t V + \mathcal{A}_t^w V + \mathcal{A}_t^z V. \tag{1.14}$$

In addition, function V satisfies a number of boundary conditions derived from the constraints on w, k, z. These state-constraint boundary conditions are detailed in Section 1.A.1.

In a model with a degenerate stationary lifecycle (i.e. infinitely lived or perpetual youth households), the coefficients of the HJB equation are all independent of t, and policy functions are all independent of age. It then suffices to look for a stationary solution $V(w,z)$, which can typically be found by guessing and checking a value function resulting from a consumption function linear in wealth.

With realistic lifecycle, however, we need to solve the full age-dependent version of the HJB equation above, for which no analytic solution typically exists. Instead, I set out in Section 1.B.1 how to find a numerical solution using a finite difference method.

Agents' policy functions satisfy simple first-order conditions, with constraints only binding on the boundary of the state space. In continuous time, first-order

conditions are static, and relate to the state-space derivatives of the value function. The usual intertemporal Euler equations also hold.

Proposition 1.1 (First order conditions). *The household's optimal policy functions $c(t,w,z)$, $q(t,w,z)$, $k(t,w,z)$ for consumption, wealth at death, and capital holdings respectively, evaluated at age t, wealth level w, and productivity z, are characterized by the first order conditions*

$$c(t,w,z) = (u')^{-1}\big(\partial_w V(t,w,z)\big) \tag{FOC c}$$

$$q(t,w,z) = (\tilde{u}')^{-1}\left(\frac{\pi_t}{\lambda_t}\partial_w V(t,w,z)\right) \tag{FOC q}$$

$$k(t,w,z) = \min\left\{-\frac{\partial_w V(t,w,z)}{\partial_{ww} V(t,w,z)}\frac{r^k - r^b_t}{\sigma^2}, w - \Phi\right\}, \tag{FOC k}$$

for any (t,w,z) such that $t > 0, w > \Phi$.

Proposition 1.2 (Euler equations). *Let (w_t, z_t) denote the optimal paths of wealth and productivity given an initial value and draws of the shock processes. Then the path of consumption $c_t = c(t,w_t,z_t)$ satisfies*

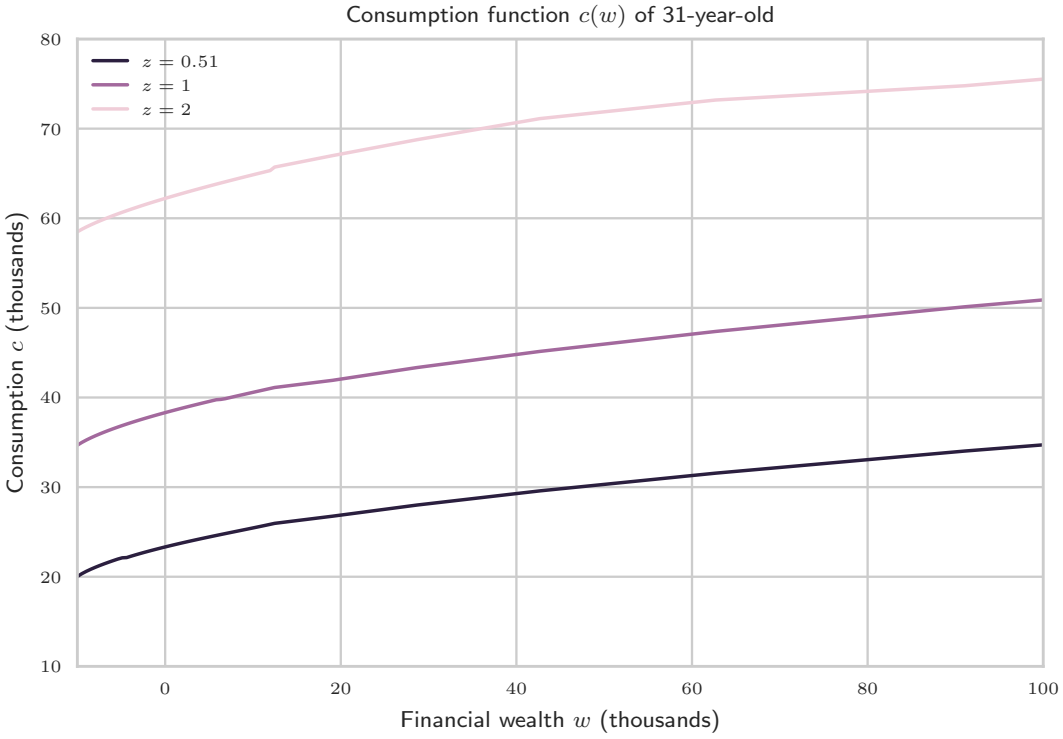
$$\begin{aligned}(\rho - r^b_t + \pi_t - \lambda_t - \partial_w \Theta(t,w_t,z_t))\,dt &= E\frac{du'(c_t)}{u'(c_t)} + \frac{dk_t}{k_t}\frac{du'(c_t)}{u'(c_t)} \\ &= E\frac{d\tilde{u}'(q_t)}{\tilde{u}'(q_t)} + \frac{dk_t}{k_t}\frac{d\tilde{u}'(q_t)}{\tilde{u}'(q_t)}\end{aligned}$$

away from the constraint, i.e. whenever $w_t > \Phi$.

The Euler equations are the continuous-time versions of the familiar expressions, generalized to age-specific mortality, annuities, and transfers. For example, $du'(c_t)/u'(c_t)$ is the growth rate of marginal utility, and $dk_t/k_t \cdot du'(c_t)/u'(c_t)$ is its covariance with the return on capital. Note that when annuity markets are competitive ($\pi_t = \lambda_t$) then agents fully smooth consumption against their mortality

FIGURE 1.3: CONCAVE CONSUMPTION FUNCTIONS

Note: Figure plots consumption functions of financial wealth for a 31 year-old at different productivity levels ($z = 1$ corresponds to the median)



risk λ_t. Asset-tested transfers (where $\partial_w \Theta \neq 0$) distort households' consumption decisions.

Additionally, we shall see later that household consumption policy functions are concave and asymptotically linear in wealth for large w. This is illustrated by Figure 1.3, which plots the consumption function out of financial wealth for a 31 year old with different productivity levels (for the calibration set out in the next section). We see relatively high marginal propensities to consume at low levels of wealth. The asymptotically linear behavior of consumption will be important

in Chapter 2 for characterizing the shape of the tail of the wealth distribution.

1.2.2 *Distribution of wealth and productivity within age groups*

Due to uncertainty in productivity and returns to capital, there is a nontrivial distribution of wealth and productivity within each age group.

After solving the household's problem, Proposition 1.1 gives us the optimal policy functions. Letting

$$s(t,w,z) = y_t z + \Theta(t,w,z) + r_t^b w + (r^k - r_t^b)k(t,w,z) - c(t,w,z) + \pi_t(q(t,w,z) - w) \tag{1.15}$$

denote the optimal savings function, write the budget constraint equation (1.7) as

$$\mathrm{d}w_t = s(t,w_t,z_t)\,\mathrm{d}t + \sigma k(t,w_t,z_t)\,\mathrm{d}B_t. \tag{1.16}$$

Equation (1.16) tells us how the wealth of any individual household will evolve. Combined with the distribution of wealth today, we can use this to predict the distribution of wealth tomorrow. This intuition can be formalized using standard results in stochastic analysis, which allow us to write a law of motion for the joint density of wealth and productivity within an age group, $g(w,z|t)$. The density g satisfies the *Kolmogorov Forward equation*

$$\begin{aligned} \partial_t g(w,z|t) = &-\lambda_t g - \partial_w[sg] + \partial_{ww}\left[\frac{(\sigma k)^2}{2} g\right] \\ &- \partial_z[\mu^z g] + \eta \int_{-\infty}^{\infty} \big(g(w,z'|t) - g(w,z|t)\big)\phi^z(z')\,\mathrm{d}z', \end{aligned} \tag{KF}$$

with boundary condition $g(w,z|0) = g_0(w,z)$ for some initial distribution g_0 to be discussed shortly. Here ϕ^z is the density function for the jumps in process

z_t. Section 1.A.2 outlines a derivation of equation (KF). Details of the numerical solution method can be found in Section 1.B.1.[8]

Thinking of $g(w, z|t)$ as a joint density of wealth and productivity conditional on age, we can multiply it by the density of ages $D(t)$ to get the joint density of wealth, productivity, and ages in the economy. Denote the associated probability measure by $m(\mathrm{d}t, \mathrm{d}w, \mathrm{d}z)$. This is the main distributional object in the model. It lets us work out household aggregates. For example, aggregate capital is given by

$$\begin{aligned} K &= \int k(t, w, z)\, m(\mathrm{d}t, \mathrm{d}w, \mathrm{d}z) \\ &= \int_{t\geq 0} \int_{w\in\mathbb{R}} \int_{z\in\mathbb{R}} k(t, w, z) g(w, z|t) D(t)\, \mathrm{d}t\, \mathrm{d}w\, \mathrm{d}z. \end{aligned}$$

Initial distribution

The joint distribution of wealth and productivity for entering agents, $g_0(w, z)$, is an important modeling choice. For simplicity, I assume initial wealth and productivity are independent, but this assumption can to be relaxed for certain applications.

Empirically, income inequality is large for young households and continues to widen across the lifecycle. To replicate this fact, I take assume initial productivity follows a censored version of the ergodic (i.e. long-term) distribution of productivity.

[8]Following Achdou et al. (2017), I exploit a deep symmetry between equations (HJB) and (KF), by which equation (KF) can be written in terms of the L^2-adjoint operator of the infinitesimal generator $\mathcal{A}_t$, namely

$$\partial_t g = \mathcal{A}_t^* g - \lambda_t g.$$

Solving (HJB) yields the operator $\mathcal{A}_t$, which makes it easy to solve (KF).

To close the model, the distribution of initial wealth should be linked appropriately to the distribution of bequests made at every point in time. The model does not keep track of family relationships between agents. Instead, "your child" is the entering agent who happens by chance to inherit your wealth when you die.

Of course, in real life, bequests are not necessarily received at the time of household formation. Instead, this could be understood as agents borrowing against their future bequest, or as the monetary value of an intangible transmission of human wealth at economic birth.

Let

$$\Lambda = \int_0^\infty \lambda_t D(t)\,\mathrm{d}t$$

denote the total mass of agents dying and e the total mass of agents entering at a given instant.

Recall that agents choose the optimal size q of their total bequest if they were to die in this instant. Let $\tilde{q} = q + \vartheta^q(q)$ denote the size of the bequest after any estate taxes and redistribution.[9] This is split into e/Λ many gifts, each of size $\tilde{q}\Lambda/e$. An entering agent is given a random draw from the distribution of split gifts from agents that actually die. The density function for initial wealth is thus

$$f_0^w(w) = \frac{e}{\Lambda}\int_0^\infty f^{\tilde{q}}\left(\tfrac{e}{\Lambda}w|t\right)\lambda(t)D(t)\,\mathrm{d}t, \tag{1.17}$$

Here $f^{\tilde{q}}(\tilde{q}|t)$ is the density of after-tax bequests $\tilde{q}$ for agents of age t, obtained using the after-tax bequest policy function $q(t,w,z) + \vartheta^q(q(t,w,z))$.

[9]Note that in the model as set out above, it is pre-tax bequests q that enter the utility function rather than $\tilde{q}$. In the exercises presented in this thesis there are no bequest taxes – it would be straightforward to incorporate them into the utility function to correct this in a relevant application.

We can see that the total mass of newborn wealth,

$$e \int w f_0(w) \, \mathrm{d}w,$$

equals the total mass of after-tax bequests made from dying agents.

1.2.3 *Government*

Let τ denote calendar time. The schedule of taxes and subsidies in place at time τ is denoted $\Theta_\tau(t, w, z)$. This can depend on age, financial wealth, and productivity.

Government consumption c^g_τ is chosen exogenously. Finally, any budget deficit or surplus is held in bonds b^g_τ. Accordingly, the government budget constraint is

$$\frac{\mathrm{d}b^g_\tau}{\mathrm{d}\tau} = r^b_\tau b^g_\tau - c^g_\tau - \int \Theta_\tau \, \mathrm{d}m_\tau - \int \vartheta^q_\tau \, \mathrm{d}m_\tau, \tag{1.18}$$

where the last two terms denote the net receipts of the transfer programs and of the estate tax, respectively.

1.2.4 *Equilibrium*

Recall that in the presence of period or cohort effects, agents' policy functions will also vary with calendar time τ (or equivalently with their date of birth $s = \tau - t$). I reintroduce this dependence into the notation in this section.

The main equilibrium condition is that bonds are in zero net supply (Huggett 1993).[10] Recalling that $w = b + k$, aggregate household bond holdings at calendar date τ can be written as

$$b^h_\tau = \int \big(w - k(t, w, z)\big) \, m_\tau(\mathrm{d}t, \mathrm{d}w, \mathrm{d}z).$$

[10] We could also close the model using investment into aggregate productive capital, and use that to back out the equilibrium rate as in Aiyagari (1994).

Since the government also participates in the bond market, the bond market clearing condition is

$$b_\tau^g + b_\tau^h = 0 \qquad \text{at all times } \tau. \tag{EQM}$$

Next, aggregate wealth should be maintained when agents die and new households enter the economy. This is summarized by the condition

$$\int \lambda_t w \, m_\tau(\mathrm{d}t, \mathrm{d}w, \mathrm{d}z) = \int_{t=0} w \, m_\tau(\mathrm{d}t, \mathrm{d}w, \mathrm{d}z) \qquad \text{at all times } \tau. \tag{1.19}$$

An equilibrium of this model is defined as a collection of policy functions for consumption and investment for agents of each cohort, together with a path for the risk-free interest rate and for government policy, so that the policy functions solve each agent's HJB equation and so that equations (1.19) and (EQM) hold.

Equilibria in which equation (1.19) hold are computationally difficult to find. The distribution of wealth of new households affects both the left and right-hand side of the equation; solving the fixed point problem associated with equation (1.19) is quite demanding. In some economic applications – notably Chapter 4 – the intergenerational transmission of wealth is not so important. As a result, it is useful to define an associated equilibrium concept in which equation (1.19) does not need to hold: an *equilibrium without intergenerational transmission.*

A *stationary equilibrium* is one in which none of the policy functions, prices, or distribution functions vary with calendar date τ (although they can of course vary with household age t). This is the appropriate stochastic steady state concept for the model, corresponding to an economy with no period or cohort effects.

Note that the market for risky capital does not clear. Rather, the supply of risky capital is unbounded and infinitely elastic. In Section 3.2, we will want to

read extra meaning into the equilibrium rate r^b, so I shall disable the risky asset k to ensure that aggregate *wealth* is in zero net supply.

Transition equilibrium

In a transition equilibrium, we must keep track of calendar time τ alongside household characteristics of age t, wealth w, and productivity z.

The value function thus gains an argument, $V(\tau, t, w, z)$, and the HJB gains a corresponding ∂_τ term:

$$(\rho + \lambda_t)V(\tau, t, w, z) = u(c) + \lambda_t \tilde{u}(q) + \partial_t V + \partial_\tau V + \mathcal{A}_t^w V + \mathcal{A}_t^z V. \tag{1.20}$$

Since there are now three household state variables and time to keep track of, we cannot directly apply the tools of Achdou et al. (2017) without great computational cost due to the high dimensionality: the memory size of the state space grows exponentially with the number of dimensions. Instead, I take the approach of slicing up the problem by cohort. We can reduce dimensionality by one by *currying*, i.e. defining a series of functions

$$V^s(t, w, z) = V(t, w, z; s + t) \quad \text{for all } t \geq 0, \tag{1.21}$$

where each one represents the value function for a household born at time s. Note that calendar time τ, agent age t and date of birth s jointly determine *two* important state variables (since $\tau = s + t$). As a result, we have that

$$\partial_\tau V^s = \partial_t V + \partial_\tau V,$$

so that substituting into equation (1.20) eliminates a term,

$$(\rho + \lambda_t)V^s(t, w, z) = u(c) + \lambda_t \tilde{u}(q) + \partial_t V^s + \mathcal{A}_t^w V^s + \mathcal{A}_t^z V^s. \tag{1.22}$$

We can do the same for the joint distribution of wealth $g(\tau, t, w, z)$ and reduce the dimensionality of the Kolmogorov Forward equation (KF).

Accordingly, a transition equilibrium requires

1. a solution $V^s(t, w, z)$ to the household problem (HJB) for each cohort with date of birth s;

2. a joint distribution m_τ of wealth, productivity and age at each time τ such that the density function $g^s(w, z|t)$ of wealth and productivity for each cohort satisfies the relevant Kolmogorov Forward equation; and

3. a path of prices (r^b_τ) at each time such that the bond market always clears,

$$\int b(t, w, z; \tau - t) m_\tau(\mathrm{d}t, \mathrm{d}w, \mathrm{d}z) + b^g_\tau = 0 \qquad \text{for all } \tau.$$

The problem of solving the HJB and KF equations for each cohort are independent and can be computed in parallel. Once this is done, the joint wealth distribution $g(\tau, t, w, z)$ can be obtained by pasting together the cohort distributions. This dimensional reduction is enough to permit problems with relatively large grid sizes to be solved quickly.

Further numerical issues that crop up in extending the canonical methods to a lifecycle setting are addressed in Section 1.B.

1.2.5 *Solution method*

In this section, I provide a bird's eye view of the numerical methods for solving the model. More details can be found in the later sections.

The main equations of the model are equations (HJB), (KF) and (EQM). Together, these define a *mean field game:* a type of system of integral and differ-

ential equations described in Lasry and Lions (2006, 2007). To solve it, I adapt the framework for numerical solution presented in Achdou et al. (2017).

Note that agent lifespans are uncertain and have unbounded support in the model. However, numerical solution requires a bounded grid on household ages. There are several approaches to modeling behavior at the upper bound on ages. I choose to assume that after a maximal age (say 150 years), the agent's problem becomes stationary (i.e. perpetual youth). To solve the agent's problem, I first solve this stationary problem for extreme old age, and then solve the lifecycle model backwards in time from there.[11]

The following sketches the solution algorithm for a transition equilibrium *without intergenerational transmission*. (I will explain later how to find the steady-state distribution of initial wealth.)

1. Make a guess of the path of interest rates (r_τ^b).

2. Given (r_τ^b), solve the agent's problem for each cohort as follows. First, solve the stationary version forward in time for agents of age 150. Next, use a fully implicit finite difference method to solve backwards in time from the stationary solution. This involves finding the solution of a high dimensional nonlinear equation at each time step, using a Newton method. (This is fast when exploiting the sparsity of the system which results from the continuous time approach.) Use this to work out the policy functions for agents of all cohorts and ages.

[11] We could also ensure that agents die for sure once they reach age 150, by setting the value function at that age to zero and then solving backwards. In practice, it does not really matter economically what you choose for end of life if the maximal age is large enough since the probability of surviving to that age vanishes. The approach of imposing terminal stationarity has the advantage of numerical stability, since policy functions in the last period of life remain well-behaved.

3. Given policy functions, solve equation (KF) for each cohort forward in time, to get the cohort-conditional joint distribution. As with Achdou et al. (2017), I exploit the adjointness of the KF and HJB equations to solve this easily (see Section 1.A.2). Integrate the age- and cohort-specific wealth density against the demographic structure function to get the joint distribution m_τ of age, productivity, and wealth at each time τ.

4. Use this distribution to calculate excess bond holdings as the difference between aggregate capital holdings and aggregate wealth (equation (EQM)) at each time τ.

5. Update the guess of (r_τ) in accordance with the signs of the excess bond holdings (a kind of tâtonnement). This can be done through a variety of methods.[12]

6. Repeat until the bond market clears.

Several results (Barles and Souganidis 1991; Lasry and Lions 2007; Achdou et al. 2014) guarantee that numerical schemes of this kind will converge to a

[12] If $\mathbf{F}(\mathbf{r})$ maps the path $\mathbf{r} = (r_\tau)$ to the vector of excess bond holdings at each τ, the simplest way to find a zero of $\mathbf{F}$ is to iterating on guesses, updating them using

$$\mathbf{r}^{\text{new}} = \mathbf{r} - \boldsymbol{\Xi}\mathbf{F}(\mathbf{r}) \tag{1.23}$$

for some matrix of hyperparameters $\boldsymbol{\Xi}$ to be chosen.

Working values for Ξ can be found using trial and error. However, ideal would be to set $\boldsymbol{\Xi} = J_{\mathbf{F}}(\mathbf{r})^{-1}$ to the inverse Jacobian matrix for $\mathbf{F}$, so that equation (1.23) implements a Newton method. Due to the size of the problem, it is not feasible exactly to find the Jacobian.

However, relative to just guessing parameters $\boldsymbol{\Xi}$, numerical performance can be much improved in two cases. The first involves approximating the Jacobian matrix using first difference approximations for the derivatives, also known as the *secant method.* Another is to use *automatic differentiation* to approximate the Jacobian. Automatic differentiation computes numerically the derivatives of function $\mathbf{F}$ by reducing $\mathbf{F}$ to a (large) sequence of elementary operations, whose derivatives are then computed and composed with a chain rule to back out J.

solution. Existence and uniqueness results exist for stationary problems, with encouraging results for the time-dependent case. Section 1.A.2 contains more discussion on this topic.

To find a steady-state equilibrium with intergenerational transmission, in which the distribution of bequests lines up with the distribution of wealth of newborns (equation (1.19)), requires an additional and computationally intensive step. Begin with a guess of the distribution of initial wealth. Without changing any model parameters, solve for a long transition from that point. As generations pass and bequests are made, the two distributions will begin to align. When equation (1.19) holds within an acceptable tolerance, stop the transition. (500 years seems to be sufficient.) The economy at that time will be in a steady-state equilibrium with intergenerational transmission.[13]

1.3 MODEL PARAMETERIZATION AND FIT

As discussed in the introduction, the aim of this chapter is to present a general form of the framework, nesting the specific models used in later chapters. These use different versions of the model and vary in their calibration strategy. In this section, I focus on a simple version used in Chapter 4, in order to illustrate how these models are calibrated and to evaluate their performance against the data. This does not have all the features of Section 1.2. Here, the bequest motive is turned off ($\chi = 0$). There is thus no intergenerational redistribution of wealth – instead, accidental bequests are wasted, and as we shall discuss, new households enter with an exogenous distribution of initial wealth. This is an *equilibrium*

[13]Some numerical tricks can speed up this process, such as solving an auxiliary version of the model with coarser lifecycle and wealth grids in order to obtain a good starting guess for the full specification.

without intergenerational transmission. One goal of this section is to see how well a lifecycle heterogeneous agent model can match the data without too many departures from the standard consumption model, using the extra computational space to calibrate these canonical mechanisms flexibly.

The time scale is annual.

1.3.1 *Household calibration*

Households are born economically at age $t = 25$. Age-specific mortality rates λ_t are taken from life tables for the U.S. in 2015, from the Human Mortality Database. Specifically, I interpolate the actuarial measure q_x, the probability of surviving from age x to $x + 1$. I assume λ_t is constant after age 90.

When solving the model numerically, we must impose an upper bound to lifespans, which I set at 150 years (agents have an ex ante probability around 10^{-12} of reaching that age in the model).

I take the age structure $D(t)$ from the U.S. Census Bureau for 2015, normalizing population size to 1 so that model aggregates correspond to per capita numbers. For the sake of simplicity, in this version I freeze $D(t)$ over time.

Preferences

The intertemporal elasticity of substitution is set to $\gamma = 1.5$, and the household discount rate ρ is calibrated internally to give an equilibrium interest rate of $r^b = 2\%$ p.a. This corresponds to $\rho = 5.6\%$.

Bequests and initial wealth

Since there are no bequests, we must pick a distribution for the initial wealth of households. I use a log-normal distribution, parameterized to match the mean ($27,000) and coefficient of variation (13.1) for households aged 25 and below in the 2016 Survey of Consumer Finances.

Assets

I calibrate r^k to 5.8% and σ to 0.2. There is no standard way of setting the parameters of the idiosyncratically risky asset k. Recall that it captures *idiosyncratic* returns on housing and private business equity. This entails looking at micro data on returns and cleaning them of aggregate fluctuations. Benhabib et al. (2011) report studies of individual return on housing and business equity in the PSID and SCF respectively, and estimate a return r^k between 6.5 and 9 percent and volatility σ around 0.20. Kartashova (2014) and Moskowitz and Vissing-Jørgensen (2002) estimate even higher returns on private equity.

An alternative approach is to use the following result. The risk premium of the risky asset determines household portfolio allocation, which affects the size of the right wealth tail. The proposition relates the average return r^k to the shape of the tail of the long-run wealth distribution:

Proposition 1.3. *Under conditions set out in Section 1.A.1, the steady state marginal density of wealth g_w is asymptotically Pareto in the right tail, i.e. there exists $\alpha > 0$ such that*

$$\lim_{w \to \infty} \frac{g_w(w)}{\alpha \cdot w^{-\zeta - 1}} = 1,$$

where the tail exponent ζ satisfies

$$\zeta = \gamma \left(\frac{2\sigma^2(\rho - r^b)}{(r^k - r^b)^2} - 1 \right). \tag{1.24}$$

This result was derived by Achdou et al. (2017) for a model with no lifecycle. I adapt their proof for this model in Section 1.A.1.

Empirically, the Pareto parameter for the tail of the U.S. wealth distribution is $\zeta \approx 1.5$. Assuming $\sigma = 0.20$, equation (1.24) implies value $r^k = 5.8\%$, slightly below estimates reported by Benhabib et al. (2011).

The borrowing constraint, which relates to short term unsecured debt, is set to -$10,000.

Few retirees buy annuities in the data (see Brown (2007) for example). Many explanations have been proposed for this long-standing puzzle, such as adverse selection (Mitchell et al. 1999), bequest motives (Friedman and Warshawsky 1990), self-insurance and liquidity requirements (De Nardi et al. 2006), and public insurance (Brown 2007).

These channels are outside the scope of the model. As a result, I allow for a small amount of life insurance by setting the return on the annuity to a fraction of the mortality rates,

$$\pi_t = 0.1 \cdot \lambda_t.$$

Model fit and elasticities are quite robust to changes in λ_t/π_t, as long as this ratio does not exceed around one-half. The reason is that mortality probabilities λ_t are very large after age 90. If these agents were to have exclusive access to an asset paying such a high return (above 10%), the wealth distribution would be counterfactually skewed in their favor.

Table 1.1: Estimated parameters of productivity process (annualized).

Parameter	Component 1	Component 2
Arrival rate	0.320	0.028
Mean reversion	0.761	0.009
SD of jumps	1.740	1.530

Source: Kaplan et al. (2018)

Income process

As discussed, log-productivity is the sum of two components,

$$\log z_t = \zeta_t^1 + \zeta_t^2,$$

following Kaplan et al. (2018). Each process ζ^i is a mean reverting process subject to jump shocks arriving at different frequencies, which pick a new productivity level from a normal distribution. Kaplan et al. (2018) estimate the shock parameters to match several moments of the distribution of income and of income changes, described by Guvenen et al. (2015). I use their estimates, reproduced in Table 1.1.[14] The first component captures frequent but short-lived variation, while the other has low-frequency persistent movements of productivity.

I set the age profile of wages y_t to the mean of pre-transfer wage income in the 2013 Survey of Consumer Finances.[15] The maximal age in the SCF data is 95, where y_t is 0. I also set y_t to 0 for agents older than 95. I smooth the empirical profile using a 5 year moving average.

The mortality and wage profiles are plotted in Figure 1.4.

[14]When solving the model numerically, I follow Kaplan et al. (2018) in using simulation to make a one-dimensional approximation to this process, in order to avoid making productivity a two-dimensional state variable.

[15]In particular, I look at question `X5702`, which asks "In total, what was your (family's) annual income from wages and salaries in 2006, before deductions for taxes and anything else".

Figure 1.4: Calibrated mortality rates and lifecycle wage

Note: Age-specific mortality λ_t taken from U.S. 2015 Human Mortality Database, and wage y_t is mean of wage income in thousands of dollars in the Survey of Consumer Finances, all annualized.



Numerical parameters

The solution algorithm requires grids for wealth and productivity with upper and lower bounds. I choose the lower bound for wealth below the wealth constraint, so that it never binds. I pick the upper bound high enough so that at least 99.99% of the probability mass of wealth lies below it. For productivity, I take the 33-point grid from the discretized process provided by Kaplan et al. (2018).

1.3.2 *Fiscal policy*

In these non-Ricardian models, with borrowing constraints, finite lifespans, and heterogeneity within and between age groups, careful attention should be paid to fiscal policy.

Taxes and subsidies to labor income

A progressive labor income tax/transfer schedule is implemented using a power law. This simple parameterization is increasingly popular in the literature due to its remarkably good fit to the U.S. tax system:[16]

$$\text{(post-fiscal earnings)} = \theta_0 \cdot \text{(pre-fiscal earnings)}^{1-\theta_1}. \tag{1.25}$$

Post-fiscal earnings are composed of market labor income $y_t z_t$ plus working-life transfers net of income taxes. (I do *not* include social security transfers, which will be described below.)

Parameter θ_1 controls the progressivity of the schedule: $\theta_1 \to 1$ corresponds to the limiting most progressive setup in which everyone gets the same after-tax income. Parameter θ_0 shifts the level; it is useful for maintaining government budget targets.

Kaymak and Poschke (2016), Heathcote et al. (2017), and Straub (2017) all estimate the progressivity parameter θ_1 in the U.S. using various methods.[17] They find a range of estimates between 0.08 and 0.2. I use the most recent estimate in the literature: Straub (2017) uses PSID data from 2013 to find a θ_1 of 0.16.

Note that the tax parameters are assumed not to depend on household age. Heathcote et al. (2019) estimate age-varying parameters for this process, and find no significant variation against age.

When the level θ_0 adjusts to keep government budget balance, it is typically around 1.5. In that case, low incomes receive a positive subsidy.

[16] See Bénabou (2000), Kindermann and Krueger (2015), Kaymak and Poschke (2016), Heathcote et al. (2017), and Straub (2017)

[17] Kaymak and Poschke (2016) target the average income tax paid by the top 1% of earners to match data from Piketty and Saez (2007), while Heathcote et al. (2017) and Straub (2017) feed PSID data into the NBER tax simulator to back out θ_1.

Table 1.2: Fit of wealth distribution

Moment	Data (SCF 2013)	Model
Gini coefficient	0.85	0.84
90/50 ratio	11.6	11.9
99/50 ratio	97	71
Mean/median ratio	6.5	5.7
Wealth/income ratio	3.1	3.4
Fraction of hand-to-mouth*	12-20%	16%

* Data is from Kaplan and Violante (2014), who define hand-to-mouth (in the net worth sense) as households with negative net worth.

Social security

Public transfers to seniors consist of a lump-sum payment to over 65's of $10,000 per year, received continuously. This simple formulation captures redistributive features of U.S. social security while abstracting from the forced saving component. This setup is sufficient for certain applications, such as Chapter 4, where I show the main results are robust to the precise specification of baseline fiscal policy.

1.3.3 *Model fit and analysis*

In this section, I examine the ability of the calibrated model to match empirical facts on inequality, on lifecycle, and on the intersection of these two axes.

Inequality and cross-section

Table 1.2 compares moments of the wealth distribution in the model and in the 2013 Survey of Consumer Finances. The model does a good job at matching the empirical Gini coefficient; however, it is missing some mass deep in the right

tail (at the top 1% and above). For this reason, it undershoots the mean-to-median ratio. The reason is twofold. Firstly, Achdou et al. (2017) and Benhabib et al. (2015) showed that income shocks (even leptokurtic) cannot alone deliver a fat Pareto right tail to the wealth distribution. The intuition is that for top asset holders, labor income is negligible compared to capital. As a result, they have a negligible precautionary savings motive and do not save enough to make the wealth tail fat.

Benhabib et al. (2011) argued that idiosyncratic risk to *capital returns* can yield a fat tailed stationary distribution of wealth. However, in Chapter 2 I argue that this mechanism is too slow in a lifecycle model to match the data on changes in inequality; it operates on the level of generations. Since here initial wealth is thin-tailed, so too is the overall distribution of wealth. Section 2.3 provides further results and discussion of the fit of the tail of the model's wealth distribution.

A fix would be to have a faster mechanism, as argued in Gabaix et al. (2016), such as superstar behavior in returns on wealth, or a positive correlation between returns and wealth (an established idea, with support from the Alstadsæter et al. (2017) data on tax evaders and explored by Benhabib et al. (2019)).

Alternatively, the initial wealth distribution could be made more leptokurtic, capturing other forms of inequality in permanent income stemming from human capital, education, and talent. Indeed, Huggett et al. (2011) and Storesletten et al. (2004) find that these are important contributors to wealth and income inequality.

Importantly, the model does well at matching the fraction of hand-to-mouth agents in the economy. These agents are defined to be those with negative net

Figure 1.5: Lifecycle profiles by productivity group

Note: Plotted in thousands of dollars. For example, if Ω is a quantiles of the marginal distribution of z_t, the top right plots $\int c(w,z,t)g(w,z|t)1_{z\in\Omega}\, m(\mathrm{d}w,\mathrm{d}z,\mathrm{d}t)$ against age for different Ω.



worth. As we shall see in Chapter 4, these agents are important for understanding the aggregate and cross-sectional response of savings to fiscal policy, and for wealth inequality.

Kaplan and Violante (2010) and Kaplan et al. (2014) stress the importance of the so-called *wealthy hand-to-mouth* agents in the data, who have large net worth but no liquid assets. They estimate that around 30% of U.S. households exhibit this behavior. I do not model wealthy hand-to-mouth behavior (which requires another asset) in this paper for simplicity and computational speed, but this could easily be added.

Lifecycle

Figure 1.5 plots average lifecycle profiles among different productivity subgroups. Consumption is hump-shaped, but flatter than income since agents partially self-insure through savings. Agents save during working life and dissave during retirement. Savings peak at the same time as income, in the mid-50s.

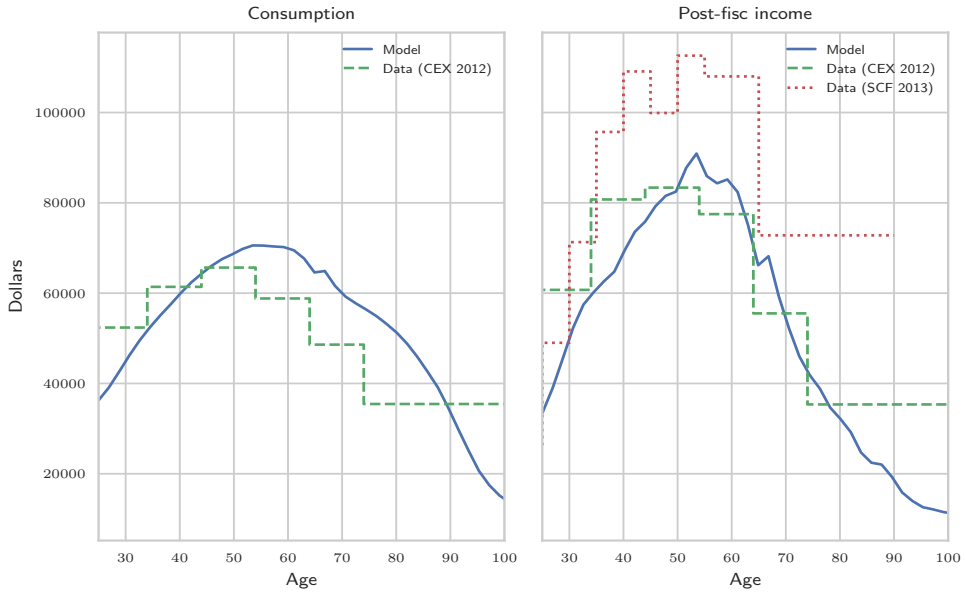
Consumption grows in early life for two reasons. Firstly, the borrowing constraint is more likely to bind among the young, who are generally less well off. As a result, they are not fully able to smooth consumption by borrowing against expected future income. Secondly, as agents age, they make their way through their permanent income. As a result, future income becomes less risky, since there is less of it. The precautionary savings motive weakens, allowing agents to increase consumption.

Consumption peaks in the late fifties and then gradually declines. This can be ascribed to uninsured mortality risk. The consumption Euler equation (Proposition 1.2) shows us that agents further discount old-age consumption by the uninsured mortality probability, pushing down the slope of consumption. As a result, consumption and wealth slowly decline to zero.

Especially important are the average profiles of consumption and income, plotted against the data in Figure 1.6. Consumption is not captured in the SCF, so I calculate profiles from data in the 2016 Consumer Expenditure Survey.

The model does matches features of the average consumption profile in the CEX during working life. There is some overconsumption in retirement – this could be ascribed to several reasons. First, a difference in data sources. The right-hand panel plots profiles of income after tax/transfers. Income is higher

Figure 1.6: Hump-shaped profiles of mean consumption and income over the lifecycle, model and data



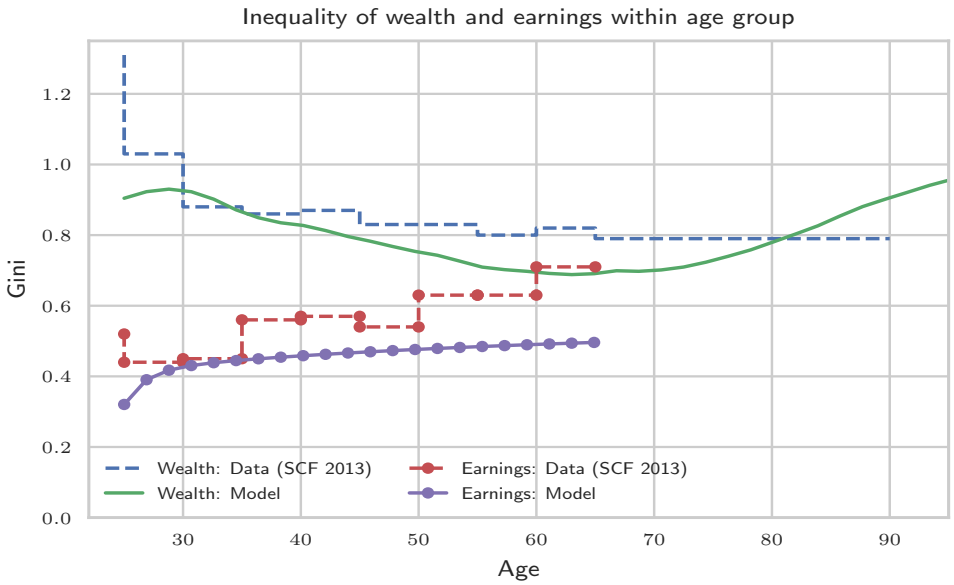
in the model than in the CEX, which could explain the overconsumption. It is still, however, below the empirical profile in the SCF. To resolve this issue would require a data source with sufficiently good quality income data to calibrate the income process while also providing data for consumption.

Second, there are cohort effects at play in the data that are not present in the steady-state equilibrium of the model. Due to growth in wages over the path centuries, current seniors will have smaller savings than the next generation, depressing their consumption.

Finally, recall that this calibration does not feature a bequest motive or health risks. These have been shown to be important savings motives, especially for seniors, which would thus lead them to reduce consumption. Adding this behavior to the model would improve lifecycle consumption profiles in an application where this is important.

FIGURE 1.7: GINI COEFFICIENTS OF WEALTH WITHIN EACH AGE GROUP

Dashed lines use 2013 data from the Survey of Consumer Finances. Data series are binned over 5 year age groups. Leftmost point aggregates ages ≤ 25 and rightmost ages > 65.



Nevertheless, it has traditionally been hard to get a hump shape in consumption that peaks in the mid 50s, as it does in the data and in the model (Thurow 1969; Bütler 2001; Fernández-Villaverde and Krueger 2007; Hansen and İmrohoroğlu 2008). The empirical literature also discusses the decline in consumption in retirement, for which see Gourinchas and Parker (2002), Feigenbaum (2008), Hurst (2008), and Olafsson and Pagel (2018).

Inequality over the lifecycle

Figure 1.7 plots Gini coefficients for wealth and income inequality by age group in the model and in the 2013 SCF, as computed by Kuhn and Rios-Rull (2016).

Income inequality gently increases over the lifecycle in the model and in the data. In the model, this is due to the compressed distribution of initial produc-

tivity diffusing over time as shocks hit. Note that the income shock calibration targets changes in income year-to-year and not the distribution. As a result, it undershoots income inequality over the lifecycle. This fit could be improved by re-estimating the shock process so that its parameters vary with age. Alternatively, the model could be modified so that households face risk to the growth rate of income.

The model has a concave profile of wealth inequality over the lifecycle that does not seem present in the data. Note that in the SCF, the Gini coefficient for households with a head aged 25 and below is very high, due to the large number of households with negative net worth. The model does not capture the full extent of inequality among these young households – there may be selection issues in the decision to form an independent household at those ages.

1.4 concluding remarks

This chapter presents a novel approach for building overlapping generations models with heterogeneity in income and wealth. The main insight is to recognize that lifecycle behavior can be treated as just another form of household heterogeneity, and thus harness recent technical advances in the study of heterogeneous agent models in continuous time, permitting a new balance between flexibility and tractability.

I present a specific version of the framework nesting canonical mechanisms of lifecycle behavior and precautionary savings. Using the extra computational space afforded by the solution method to calibrate these mechanisms carefully, I show that the model delivers empirically realistic distributions of wealth together

with moments on lifecycle consumption and savings.

The aim of this chapter is to provide a user's guide for a macroeconomist interested in these methods. The next chapters showcase the application of this framework to a number of different contexts, each using variants of the model presented in this chapter.

The work of this chapter could be extended in a number of directions depending on the application. Firstly, the production sector in this model is quite stylized by design, in order to focus on households. Firms and labor choice can be introduced for different applications, with the algorithm easily adapted to search instead for equilibrium levels of aggregate capital. For example, in Ascari et al. (2018), we join the households of this chapter to a New Keynesian production sector.

Secondly, some applications may require carefully modelling lifecycle behavior. In such settings, enhanced focus on health risk is likely to improve the fit of lifecycle moments, especially for seniors. A different shock process for labor income could also be used: either one whose parameters depend on household age, or switching entirely to a restricted income profile (RIP). The solution method can be straightforwardly modified to incorporate these changes.

Third, fiscal policy can be detailed further, to almost arbitrary granularity. For example, a debate centers on the determinants of the increase in wealth inequality in the United States. In forthcoming work, I feed the evolution of the U.S. tax and transfer system over the postwar period into the framework of this chapter, with the aim of pinning down the relative contribution of each policy by means of counterfactual experiments.

Finally, the tools supplied in this chapter permit the analysis of idiosyncratic

risk only, and focus on the transition dynamics from a sudden zero-probability change – a so-called *MIT shock.* Recent technical advances (Algan et al. 2014; Winberry 2016; Ahn et al. 2018) in the study of heterogeneous agent economies point the way to studying models that feature *aggregate risk.* Adapting these breakthroughs to the overlapping generations setting of this chapter is an important extension, allowing for the prospect of dynamic stochastic general equilibrium models featuring overlapping generations and heterogeneous agents. These will likely form a crucial part of the macroeconomist's toolkit in the future, opening the door to the many questions raised by population aging and changing distributions.

1.A Solution method: derivations and proofs

In this section, I present further derivations and proofs of the model, omitted from the main text.

1.A.1 *HJB equation*

Recall the agents' optimization problem set out in Section 1.2.1. The agents solve for their value function, defined as

$$V(t,w,z) = \max_{(c_s,k_s,z_s)_{s\geq t}} \mathbb{E}_t \int_t^\infty e^{-\rho(s-t)-\int_t^s \lambda_\xi \,\mathrm{d}\xi} \big(u(c_s) + \lambda_s \tilde{u}(q_s)\big)\,\mathrm{d}s, \tag{1.26}$$

subject to the budget constraint and borrowing limit at each age as described in the text.

Proposition 1.4 (HJB). *The value function $V(t, w, z)$ obeys the Hamilton–Jacobi–Bellman equation,*

$$\begin{aligned}(\rho + \lambda_t)V(t,w,z) = \max_{c,k,q} u(c) + \lambda_t \tilde{u}(q) + \partial_t V \\ + \partial_w V\big[y_t z + \Theta(t,w,z) + r_t^b w + (r^k - r_t^b)k - c - \pi_t(q-w)\big] \\ + \partial_{ww} V \frac{\sigma^2}{2} k^2 \\ + \partial_z V \mu^z(z) + \eta \int_{-\infty}^{\infty} (V(t,w,z') - V(t,w,z))\phi^z(z')\,\mathrm{d}z',\end{aligned} \tag{HJB}$$

on the interior of its domain, with boundary conditions that will be outlined later in this appendix.

Essentially, the HJB equation is the continuous time equivalent of the dynamic programming principle. By rewriting the optimal control problem recursively, it provides a characterization of the solution. The proof is standard in the literature.[18]

Note that the first-order conditions pin down the choice variables in the HJB equation only away from the constraint Φ. From this we understand two facts. Firstly, the HJB equation holds only for $w > \Phi$, and in fact the wealth constraint does not appear at all in the form of equation (HJB). Agents act as if they were unconstrained inside the state space. Secondly, we need a special boundary constraint to describe the behavior of V at the wealth constraint. I will discuss this condition further down.

[18] See, for example Dixit (1993) and Bardi and Capuzzo-Dolcetta (2008) for a discussion. For a proof incorporating state constraints, see Soner (1986) and Capuzzo–Dolcetta and Lions (1990). It is important to note that the right solution concept is the *viscosity solution*, for which the value function is the unique solution of the HJB equation (see Achdou et al. (2017))

Proposition 1.1 characterized policy functions in terms of the value function, derived by taking first-order conditions of equation (HJB). Note that the first-order condition for k is the general solution for a Merton (1969) style portfolio choice problem, with the addition of the borrowing constraint limiting leverage.

Proposition 1.2 set out the Euler equations for consumption and bequests. A proof is as follows.

Proof. Consider the maximized current-value Hamiltonian associated with the agent's problem, with co-states μ, ν :

$$H(w,z,t;\mu,\nu) = \max_{c,k,q} u(c) + \lambda_t \tilde{u}(q) + \nu \left[\frac{\sigma^2 k^2}{2}\right] + \mu\big[y_t z + \Theta(t,w,z) + r_t^b w + (r^k - r_t^b)k - c - \pi_t(q-w)\big]. \tag{1.27}$$

Optimality implies the following equations:

$$0 = \partial_c H = \partial_q H = \partial_k H, \quad d\mu_t = -\partial_w H \,\mathrm{d}t + \nu_t \,\mathrm{d}B_t.$$

Together with the budget constraint, these yield the result. □

HJB boundary conditions

The value function $V(t,w,z)$ has domain

$$\Omega = \left\{(t,w,z) \in \mathbb{R}^3 \mid t \geq 0, w \geq \Phi\right\}.$$

As discussed above, equation (HJB) is a partial differential equation for V that holds on the interior of Ω, while special conditions hold on the boundary of Ω.

The borrowing constraints do not appear in equation (HJB). Instead, they pin down the boundary conditions for V on $\partial\Omega$. These are *state-constraint boundary*

conditions, as in Achdou et al. (2017), Soner (1986), and Capuzzo–Dolcetta and Lions (1990).

In the presence of a general utility function for bequests, it may not be possible to find a closed form for the boundary condition satisfied by V on $\partial\Omega$. However, let us focus on the special non-homothetic case where $u(c) = c^{1-\gamma}/(1-\gamma)$ is CRRA and $\tilde{u}(q) = \chi u(q)$ for some constant $\chi \geq 0$. (This is the case studied by Benhabib et al. (2016), where parameter χ represents the relative strength of the bequest motive.)

Proposition 1.5. *Suppose* $u(c) = c^{1-\gamma}/(1-\gamma)$ *is CRRA and* $\tilde{u}(q) = \chi u(q)$. *Then* $V(t, w, z)$ *satisfies the following* state-constraint boundary condition *at the lower wealth boundary* Φ*:*

$$\partial_w V(t, \Phi, z) \geq u'\left(\frac{(r_t + \lambda_t)\Phi + y_t z + \Theta(t, \Phi, z)}{1 + \pi_t^{1-1/\gamma}(\chi\lambda_t)^{1/\gamma}}\right) \quad \textit{for all } z, t. \tag{1.28}$$

Proof. (The following calculations hold even without assuming CRRA.) Recall the current-value maximized Hamiltonian H from equation (1.27) associated to the agent's problem, with co-state $\mu = u'(c) = \partial_w V$. The optimality conditions imply that $\partial_\mu H|_{\mu=\partial_w V}$ equals the optimal drift of wealth. Capuzzo–Dolcetta and Lions (1990) show that the appropriate boundary condition on the HJB equation corresponding to the constraint on the state variable is

$$\partial_\mu H|_{\mu=\partial_w V(t,w,z)} = 0 \text{ on } \partial\Omega.$$

In other words, the optimal drift of wealth on the boundary must be equal to zero, which is intuitive for a reflecting boundary.[19] Further, we know that since

[19] See, for example, Dixit (1993).

$k \geq 0$ and $k \leq w_t - \Phi$, we must have that $k = 0$ on the lower wealth boundary where $w = \Phi$. Together, these imply

$$y_t z + r_t^b \Phi + \Theta(t, \Phi, z) - c - \pi_t(q - \Phi) = 0, \tag{1.29}$$

where c and q are optimally chosen from their first-order conditions set out in Proposition 1.1. In other words, optimal savings must be zero at the lower boundary.

Substituting in the first-order conditions, we can rewrite equation (1.29) in terms of the value function:

$$y_t z + (r_t^b + \pi_t)\Phi + \Theta(t, \Phi, z) = (u')^{-1}\big(\partial_w V(t, w, z)\big) + \pi_t(\tilde{u}')^{-1}\left(\frac{\pi_t}{\lambda_t}\partial_w V(t, w, z)\right). \tag{1.30}$$

This is the general form of the state-constraint boundary condition for our problem, and there is generally no way to rewrite equation (1.30) as a simple condition with V on one side.

If we now assume the CRRA functional forms of the proposition, the right-hand side simplifies to equation (1.28) as required. □

The intuition is simple: the state-constraint boundary condition ensures that at the lower wealth boundary, the optimal drift on wealth is nonnegative, thus pushing the agents back away from the boundary. It does this by ensuring that both consumption and spending on the annuity are low enough for the household to grow its wealth back away from the borrowing constraint. The wealth boundary is therefore a reflecting boundary.

Upper wealth constraint

When we solve the problem numerically, we must do so on a bounded grid. This implies that we need to set an upper boundary for wealth, $w \leq \bar{w}$. Best is to choose $\bar{w}$ high enough for the wealth distribution to have negligible mass above $\bar{w}$. This way, the upper bound essentially never binds in equilibrium.

This bound is another state constraint, so we need the relevant boundary condition on the HJB equation. This condition guarantees that optimal drift of wealth at the upper boundary is negative. The optimal drift (i.e. savings) must satisfy

$$y_t z + \Theta(t, \bar{w}, z) + r_t^b \bar{w} + (r^k - r_t^b) k(t, \bar{w}, z) - c(t, \bar{w}, z) - \pi_t (q(t, \bar{w}, z) - w) \leq 0. \quad (1.31)$$

We need to translate this into a boundary condition for $V(t, w, z)$. To do so, we need to have an idea of the behavior of k for large w. In the CRRA case, I show below that the value function is asymptotically a power function with curvature $1 - \gamma$ in the wealth tail.[20] Then for large enough $\bar{w}$, risky capital holdings satisfy

$$\begin{aligned} k(t, \bar{w}, z) &\approx -\frac{\partial_w V(t, \bar{w}, z)}{\partial_{ww} V(t, \bar{w}, z)} \frac{r^k - r_t^b}{\sigma^2} \\ &= \frac{r^k - r_t^b}{\gamma \sigma^2} \bar{w}. \end{aligned}$$

Substituting this into equation (1.31), we obtain the condition

$$c(t, \bar{w}, z) + \pi_t q(t, \bar{w}, z) \leq y_t z + \Theta(t, \bar{w}, z) + \bar{w} \left(r_t^b + \lambda_t + \frac{(r^k - r_t^b)^2}{\gamma \sigma^2} \right).$$

When $\tilde{u} = \chi u$, as we assumed when discussing the lower boundary condition, we can use the first-order conditions for c and z to get the boundary condition

[20] I follow Achdou et al. (2017, Online appendix, pp15–6) here.

at the upper boundary $\bar{w}$:

$$\partial_w V(t,\bar{w},z) \geq u'\left(\frac{y_t z + \Theta(t,\bar{w},z) + \bar{w}\left(r_t^b + \lambda_t + \frac{(r^k - r_t^b)^2}{\gamma\sigma^2}\right)}{1 + \pi_t^{1-1/\gamma}(\chi\lambda_t)^{1/\gamma}}\right) \quad \text{for all } t, z. \tag{1.32}$$

Asymptotic linearity

Above and in Section 1.3.1, I argued that household policy functions are asymptotically linear in wealth, and that as a result it is possible to back out an expression for the shape of the tail of the wealth distribution. I provide a sketch proof below. More details – in the case of a non-lifecycle model – can be found in Achdou et al. (2017) and Benhabib et al. (2016).

Lemma 1.1. *When the utility function is CRRA as in Proposition 1.5, the agent's policy functions are asymptotically linear in wealth as $w \to \infty$. Namely, for any t and z,*

$$\frac{c(t,w,z)}{w} \to \zeta_t^c := \frac{\rho + \lambda_t - \frac{1}{2}\frac{(r^k - r_t^b)^2}{\sigma^2}\frac{1-\gamma}{\gamma} - (1-\gamma)(r_t^b + \lambda_t)}{\gamma\left(1 + \lambda_t\chi^{1/\gamma}\right)}, \tag{1.33}$$

$$\frac{q(t,w,z)}{w} \to \zeta_t^q := \frac{\rho + \lambda_t - \frac{1}{2}\frac{(r^k - r_t^b)^2}{\sigma^2}\frac{1-\gamma}{\gamma} - (1-\gamma)(r_t^b + \lambda_t)}{\gamma\left(\chi^{-1/\gamma} + \lambda_t\right)}, \tag{1.34}$$

$$\frac{k(t,w,z)}{w} \to \zeta_t^k := \frac{r^k - r_t^b}{\gamma\sigma^2}, \tag{1.35}$$

$$\frac{s(t,w,z)}{w} \to \zeta_t^s := \frac{r_t^b - \rho}{\gamma} + \frac{(r^k - r_t^b)^2}{\gamma\sigma^2}\frac{1+\gamma}{2\gamma}. \tag{1.36}$$

Note that the asymptotic slopes $\zeta_t^c, \zeta_t^q, \zeta_t^k, \zeta_t^s$ can vary with age.

Proof sketch. Consider an auxiliary version of the household problem, with no borrowing constraint, income, or mortality risk. Then guessing and checking a value function proportional to $w^{1-\gamma}$ yields the solution to the corresponding HJB

equation. In this problem, Equations (1.33) to (1.36) can be shown to hold with equality.

Using a limiting argument, in which borrowing constraint and wealth are gradually introduced, we can see that the solution to the full model converges to that of the auxiliary problem as wealth gets large.

The intuition is that when wealth is large enough, the distortions to the consumption–savings problem are relatively so small as to be negligible compared to the return on wealth. Hence, the solution to the real HJB equation is very close to that of the auxiliary problem outlined above. □

If policy functions are linear for large w, this means that the distribution of wealth behaves like a reflected geometric Brownian motion in the tail. Standard results (Champernowne 1953; Gabaix 2009) show that the stationary distribution of such a process is asymptotically Pareto in the tail, and let us back out the tail exponent ζ, proving Proposition 1.3.

1.A.2 *Wealth distribution*

The evolution of the joint distribution of wealth and productivity for a given cohort is pinned down by the *Kolmogorov Forward* (also known as *Fokker–Planck*) equation. This result derives a partial differential equation for the distribution of a stochastic process, given its dynamics. Please see Karatzas and Shreve (1991) for a proof.

Theorem 1.1 (Kolmogorov, Fokker, Planck). *Let w_t be a random process obeying the SDE*

$$\mathrm{d}w_t = \mu(w_t, t)\,\mathrm{d}t + \sigma(w_t, t)\,\mathrm{d}B_t$$

for a Brownian motion B_t. If $f(w,t)$ denotes the density function of w_t then f obeys the Kolmogorov forward equation

$$\partial_t f(w,t) = -\partial_w(f(w,t)\mu(w,t)) + \partial_{ww}\left(\frac{\sigma^2(w,t)}{2}f(w,t)\right). \tag{1.37}$$

We can define a differential operator $\mathcal{L}$ as the right-hand side of equation (1.37),

$$\mathcal{L}f(w,t) = -\partial_w(f(w,t)\mu(w,t)) + \partial_{ww}\left(\frac{\sigma^2(w,t)}{2}f(w,t)\right).$$

It turns out that $\mathcal{L}$ is the adjoint operator[21] to the *infinitesimal generator* $\mathcal{A}$ of the process w. The infinitesimal generator is defined by

$$\mathcal{A}f(w,t) = \mu(w,t)\partial_w f + \frac{\sigma^2(w,t)}{2}\partial_{ww} f, \tag{1.38}$$

and is an important differential operator in stochastic analysis which shows up often. Indeed, a more general formulation of the Kolmogorov Forward equation states that for any stochastic process w_t with generator $\mathcal{A}$, the density function $g(w,t)$ satisfies

$$\partial_t g = \mathcal{A}^* g.$$

The generator $\mathcal{A}$ appears on the right-hand side of equation (HJB). This deep symmetry linking the HJB and KF equations, where their underlying operators are adjoint, turns out to have wide-reaching consequences. Following Achdou et al. (2017), we can exploit this relationship to obtain a numerical solution to the KF equation for free after solving the HJB equation. See footnote 27 of Section 1.B.1 for details.

[21] These are adjoints in the L^2 sense, meaning that $\int(\mathcal{L}f)\cdot g = \int f\cdot(\mathcal{A}g)$ for any functions f and g.

Applying Theorem 1.1 to the dynamics of wealth given in equation (1.16) gives a KF equation for the wealth of an individual, conditional on survival.[22] However, probability mass at any age is lost by agents dying. As a result, we must subtract off this mass $\lambda(t)g(w, z|t)$ from the right-hand side; thus yielding equation (KF) for the density of wealth held at each age for a given cohort.

Existence and uniqueness of equilibrium

Existence and uniqueness results for mean field games, such as the one in this paper, have not yet been fully established, but promising results are forthcoming. Existence and uniqueness has been demonstrated for stationary solutions. For time-dependent solutions, however, we only have existence results so far. Please see Achdou et al. (2017) and Lasry and Lions (2007, 2006) for further discussion.

Anecdotally, the algorithm always converges to the same thing from different initial guesses, which is encouraging. Indeed, numerical experiments indicate that excess savings are a monotonic function of the interest rate.

If we consider the subproblem of solving the HJB and KF equations given interest rates, there exist many existence and uniqueness results for the solution of these PDEs. See, for example Bardi and Capuzzo-Dolcetta (2008). Note that smooth solutions to these systems do not necessarily exist. Instead, the correct solution concept – for which existence and uniqueness is guaranteed – is the *viscosity solution* of the HJB equation and the *measure-valued solution* of the KF equation. Achdou et al. (2017) provide a useful description.

[22]Strictly speaking, we need to apply an extension of Theorem 1.1 to a jump-diffusion process. See Karatzas and Shreve (1991) for these results.

1.A.3 *Demographic dynamics*

Recall from Section 1.2 that $D(\tau, t)$ denotes the density function of agents of age t at time τ. As discussed previously, at any time τ there is a measure

$$\int_0^\infty \lambda(t,\tau)D(t,\tau)\,\mathrm{d}t$$

of agents dying, together with a mass

$$\int_0^\infty e(t,\tau)D(t,\tau)\,\mathrm{d}t$$

of agents entering.

Proposition 1.6 (Demographic dynamics). *Given entry scheme e and death probabilities λ, together with an initial demographic distribution $D(t,0)$, the demographic distribution $D(t,\tau)$ at any time satisfies the PDE*

$$\partial_\tau D(t,\tau) = -\partial_t D(t,\tau) - \lambda(t,\tau)D(t,\tau) \tag{1.39}$$

on $(t,\tau) \in (0,\infty)^2$, with boundary conditions for all t, τ:

$$D_0(t) = \int_0^\infty e(t,\tau)D(t,\tau)\,\mathrm{d}t \tag{1.40}$$

and

$$D(t,0) = D_0(t). \tag{1.41}$$

Proof. As argued in Section 1.2, when $t > 0$ and $\tau > 0$, the mass of agents alive of age t at time $\tau + \delta$ equals the mass of agents of age $t - \delta$ at time τ minus those that did not survive to time $\tau + \delta$, when δ is small enough. Hence we have

$$D(t, \tau+\delta) \approx D(t-\delta,\tau)(1-\lambda(t-\delta)\delta).$$

Subtracting $D(t,\tau)$ from both sides, dividing by δ and taking the limit as $\delta \to 0$ yields the result. Additionally, the mass of agents of age zero at any time are those that have just entered, giving us the boundary condition. □

Proposition 1.6 fully characterizes the dynamics of the age structure.[23] This system can be solved numerically, for example using a finite difference method forward in time.

Analytic solutions exist only in very few cases. For example, in a perpetual youth model with $\lambda(t,\tau) \equiv \lambda$ and $e(t,\tau) \equiv \lambda$ (perfect replacement), the stationary solution can be shown by guessing and checking to be

$$D(t) = \lambda e^{-\lambda t}.$$

1.B SOLUTION METHOD: NUMERICAL STRATEGY

1.B.1 *Household problem*

Recall the household's problem set out in Section 1.2.1 and the associated discussion of the PDEs and boundary conditions in Section 1.A. In this section, I outline the numerical methods used for solving for different types of equilibria of the model.

HJB equation

With lifecycle behavior, we need to solve the full time-dependent version of the HJB equation. Although the age domain is $[0,\infty)$, numerical solution requires the imposition of an upper bound T on agents' ages. The best way to implement this is to choose a large value of T and assume that from then on, households

[23] Note that we also need that $D_0(0) = \int_0^\infty e(t,\tau) D_0(t)\,\mathrm{d}t$ for regularity.

solve a fully stationary problem where none of the parameters depend on age. In other words, after age T the agent's policy function no longer changes and there are no further lifecycle effects – perpetual youth in old age.

I refer to this problem as the household's problem at infinity. The idea is to begin by solving (forward in time) for the value function at infinity. I then impose this as a terminal condition (at age T) for the lifecycle household's problem the value function of the agent at age T, and solve backwards.

Following Achdou et al. (2017), I solve the HJB equation numerically on a grid using an upwind finite-difference method. This is a scheme in which derivatives are replaced with their numerical approximations at each grid point, with the algorithm determining implicitly when to use forward or backwards difference approximations. See Candler (1998) and the appendices to Achdou et al. (2017) for more details. However, while these studies use a semi-implicit numerical method, we must instead use a fully implicit scheme, as I shall explain below. This family of finite difference methods (whether explicit, semi-implicit or fully implicit) was shown by Barles and Souganidis (1991) to converge to the unique viscosity solution of the HJB equation.[24]

I follow the approach from the numerical appendix to Achdou et al. (2017). Suppose that $i = 1, \ldots, I$ indexes the wealth grid $\underline{w} = w_1 < w_2 < \cdots < w_I = \bar{w}$, that $j = 1, \ldots, J$ indexes productivity $z_1 < \cdots < z_J$, and that $n = 1, \ldots, N$ indexes age on the grid $0 = t_1 < t_2 < \cdots < t_N = T$. Note that the grids do

[24]The viscosity solution is a generalized kind of solution for partial differential equations. HJB equations do not always have a smooth solution – instead it is necessary to consider functions with kinks; a feature which viscosity solutions allow for. HJB equations have been proven under rather general conditions to have a unique viscosity solution, making this the appropriate solution concept for these problems. See for example Bardi and Capuzzo-Dolcetta (2008) for details or Achdou et al. (2017) for a good overview.

not need to be evenly spaced – indeed, nonlinear grids are very useful for putting enough computational weight on the long wealth tail.

Recall the Hamilton–Jacobi–Bellman equation, equation (HJB):

$$\begin{aligned}(\rho+\lambda_t)V(t,w,z) = \max_{c,k,q} u(c)+\lambda_t\tilde{u}(q)+\partial_t V \\ +\partial_w V\big[y_t z+\Theta(t,w,z)+r_t^b w+(r^k-r_t^b)k-c-\pi_t(q-w)\big] \\ +\partial_{ww}V\frac{\sigma^2}{2}k^2 \\ +\partial_z V\mu^z(z)+\eta\int_{-\infty}^{\infty}(V(t,w,z')-V(t,w,z))\phi^z(z')\,\mathrm{d}z'\end{aligned}$$

By discretizing it onto the grids, we get the following equation for the discretized value V_{ij}^n of wealth w_i and productivity z_j at age t_n:

$$\begin{aligned}(\rho+\lambda^n)V_{ij}^n = u(c_{ij}^n)+\lambda^n\tilde{u}(q_{ij}^n)+\frac{V_{ij}^{n+1}-V_{ij}^n}{\Delta t^n} \\ +\partial_w V_{ij}^n\big[y^n z_j+\Theta_{ij}^n+r^{b,n}w_i+(r^k-r^{b,n})k_{ij}^n-c_{ij}^n-\pi^n(q_{ij}^n-w_i)\big] \\ +\partial_{ww}V_{ij}^n\frac{\sigma^2}{2}(k_{ij}^n)^2 \\ +\partial_z V_{ij}^n\mu^z(z_j)+\eta\sum_{j'}(V_{ij'}^n-V_{ij}^n)\phi^z(z_{j'})\end{aligned} \tag{1.42}$$

where $\partial_w V_{ij}^n$ and $\partial_z V_{ij}^n$ are upwind difference approximation[25] in w and $\partial_{ww}V_i^n$ is a second forward difference. The choice variables $c_{ij}^n, q_{ij}^n, k_{ij}^n$ are determined from

[25]The upwind approximation equals the forward difference $\partial_w^F V_{ij}^n = (V_{i+1,j}^n - V_{ij}^n)/\Delta w_i$ when the drift on wealth (i.e. optimal savings) is positive and the backward difference $\partial_w^B V_{ij}^n = (V_{ij}^n - V_{i-1,j}^n)/\Delta w_i$ when savings are negative.

More precisely, the savings policy function is computed from the approximated policies according to

$$s_{ij}^n = y^n z_j + r^{b,n}w_i + (r^k - r^{b,n})k_{ij}^n - c_{ij}^n - \pi^n(q_{ij}^n - w_i).$$

We can define the forward approximation for the consumption policy as $(c_{ij}^n)^F = (u')^{-1}(\partial_w^F V_{ij}^n)$ as well as the backwards approximation $(c_i^n)^B$, using the same difference approximation of the value function. By doing the same thing for the other choice variables, we can compute the forward and backwards savings rates $(s_{ij}^n)^F$ and $(s_{ij}^n)^B$. The upwind approximation to V is

the numerical derivatives of V using the first-order conditions (Proposition 1.1). For example, we have

$$c_{ij}^n = (u')^{-1}(\partial_w V_{ij}^n).$$

The final sum term of equation (1.42) corresponds to a discretization of the jump process to productivity. In practice, we can avoid doing this by hand, and simply supply the Markov transition matrix $\mathbf{A}^z$ for a shock process defined on the space of productivity grid points. We then replace the infinitesimal generator $\mathcal{A}^z$ with matrix $\mathbf{A}^z$ in the discretization.

We will be solving the HJB equation recursively, backwards in time (age) from a terminal condition. This means that given the matrix V^{n+1} we want to solve equation (1.42) for V^n. Note, however, that this is a highly nonlinear system in V^n, as all the choice variables $c_{ij}^n, q_{ij}^n, k_{ij}^n$ all depend on V^n itself. For this reason, equation (1.42) is called a *fully implicit* finite difference scheme.

In practice, if policy functions do not change too much from one period to

---

then defined as

$$\partial_w V_{ij}^n = \begin{cases} \partial_w^F V_{ij}^n & \text{when } (s_{ij}^n)^F > 0; \\ \partial_w^B V_{ij}^n & \text{when } (s_{ij}^n)^B < 0; \\ \overline{\partial_w} V_{ij}^n & \text{otherwise,} \end{cases}$$

where $\overline{\partial_w} V_i^n$ is set to force savings equal to zero. (If there were only one choice variable, consumption, then it would be set to the marginal utility of income, thus forcing the agent to consume all income.)

Candler (1998) explains why the upwind method is necessary for the convergence of the finite difference scheme.

the next, it is simpler to approximate equation (1.42) using the equation

$$\begin{aligned}(\rho+\lambda^n)V_i^n =& u(c_{ij}^{n+1}) + \lambda^n \tilde{u}(q_{ij}^{n+1}) + \frac{V_{ij}^{n+1}-V_{ij}^n}{\Delta t^n} \\ &+ \partial_w V_{ij}^n \left[y^n z_j + r^{b,n} w_i + (r^k - r^{b,n})k_{ij}^{n+1} - c_{ij}^{n+1} - \pi^n (q_{ij}^{n+1} - w_i)\right] \\ &+ \partial_{ww} V_{ij}^n \frac{\sigma^2}{2}(k_{ij}^{n+1})^2 \\ &+ \partial_z V_{ij}^n \mu^z(z_j) + \eta \sum_{j'} (V_{ij'}^n - V_{ij}^n)\phi^z(z_{j'})\end{aligned} \tag{1.43}$$

Here we use policy functions derived from *tomorrow's* (known) value function instead of today's. This kind of scheme is referred to as *semi-implicit*, and this is the method outlined by Achdou et al. (2017). It has the advantage that, given V^{n+1}, today's value V_i^n is the solution of a linear problem. Further, by substituting in the definition of the upwind derivatives into the equation, the problem reduces to a sparse matrix inversion problem, which can be solved very efficiently. This comes from the fact that difference terms depend only neighboring points on the wealth grid, $V_{i\pm1}^n$, since derivatives are inherently local.

The limitation of this approach, however, is that it uses c^{n+1} instead of c^n when searching for the value function at time t_n. In a relatively stationary problem this is a mild approximation, however in a lifecycle model there can be large differences between the consumption policy functions of today and tomorrow, especially near the borrowing constraint, and especially if the constraint is moving over time.

For example, if income is zero after a certain fixed retirement age t_r then the natural wealth constraint from that age onward is zero, and the value of negative wealth in retirement is negative infinity, since those agents will never be able to consume. Hence, the consumption policy function for negative wealth at age t_r

is not defined.

However, one time step before t_r, the agent will receive an income and is therefore able to borrow against it and obtain slightly negative wealth while remaining solvent. Hence an agent at time $t_r - \Delta t$ with slightly negative wealth will be able to consume positively. Consumption at time t_r is therefore very different from the policy function at time $t_r - \Delta t$.

It turns out that for our lifecycle model, the semi-implicit approach generally does not work. Instead, we need to solve the fully implicit nonlinear system, equation (1.42), for which the sparse matrix methods do not directly apply.

I solve this nonlinear problem numerically using a Newton method to find V^n. It turns out that the Newton method is equivalent to starting from some initial guess and using a semi-implicit method to solve forward *within each time-step.* [26]

Because this involves nested loops, in which we iterate at each time-step, notation quickly becomes unwieldy, so I proceed with a simplified example. Suppose the fully implicit discretized HJB equation we wish to solve is

$$\rho V_i^n = u(c_i^n) + \partial_w V_i^n s_i^n + \frac{V_i^{n+1} - V_i^n}{\Delta t}.$$

As discussed earlier, this is a nonlinear equation in V^n, where savings, consumption and the upwind derivatives all depend nonlinearly on forward and backward differences of V^n.

Let us change notation, and denote by V^* the value tomorrow (previously V^{n+1}) – known, since we are solving backwards – and simply denote by V the

[26] I am grateful to Benjamin Moll for his comments on this.

value for today which we are trying to find. Then

$$\rho V_i = u(c_i) + \partial_w V_i s_i^n + \frac{V_i^* - V_i}{\Delta t}. \tag{1.44}$$

The Newton method solves this equation for V by iterating forward over m from some initial guess V^0, using the update step

$$\rho V_i^{m+1} = u(c_i^m) + \partial_w V_i^m s_i^m + \frac{V_i^* - V_i^{m+1}}{\Delta t}. \tag{1.45}$$

It is important to note that the last term is an approximation of the time derivative using tomorrow's value V^* and not the next guess of today's value V^{m+1}. When this scheme converges and $V^m = V^{m+1}$, we will have obtained a function that satisfies the fully implicit equation (1.44).

Equation (1.45) is isomorphic to the update step in a semi-implicit upwind finite difference method. These can be solved efficiently, following Achdou et al. (2017). If we denote by $V^m = (V_i^m)_i$ the I-dimensional vectorization of the value function, we can rewrite equation (1.45) as the I-dimensional equation

$$\rho V^{m+1} = u(c^m) + \mathbf{A}_m V^{m+1} + \frac{V_i^* - V_i^{m+1}}{\Delta t} \tag{1.46}$$

for an appropriate matrix $\mathbf{A}_m$. This problem has a recursive structure in which V_i^{n+1} depends only on neighboring grid points $V_{i-1}^n, V_i^n, V_{i+1}^n$. This comes from the continuous formulation of the problem, and ensures that in the discretization of the problem we need only look one grid point in any direction. This recursive structure ensures that $\mathbf{A}_m$ is a highly sparse matrix, so equation (1.46) can be solved efficiently using sparse matrix inversion methods, which are implemented in most scientific computing packages.

When the scheme converges at iteration M, the matrix $\mathbf{A} = \mathbf{A}_M$ can be interpreted as the transition matrix for the value function, from t^n to t^{n+1}. This

provides an interpretation of the solution scheme as reducing the HJB equation to the form of a standard discrete-time Bellman equation associated to a discrete optimal control problem.[27]

To sum up, in the fully implicit method, we iterate the HJB equation backwards over time. At each time step, we solve for today's value given tomorrow's using a Newton method that amounts to running a semi-implicit method *forward* in equation (1.45). In practice, this Newton method is relatively fast, only requiring three or so iterations.

Boundary conditions

At the boundary of the wealth grid, certain numerical differences do not exist. For example, the backward difference $\partial_w^B V_{ij}^n = (V_{ij}^n - V_{i-1,j}^n)/\Delta w_i$ is not defined when $i = 1$. Instead, for those points, we define the missing derivatives by applying the state-constraint boundary conditions from Proposition 1.5 with equality. It turns out that the upwind derivative always picks the correct difference for this to work (see footnote 25).

Special care needs to be taken when the state constraints are moving over time, which can happen with the natural borrowing constraint, for example. Normally, the state constraint ensures that consumption is low enough for savings to be zero (or barely positive) at the boundary. However, if the boundary condition at

[27] In fact, standard results from control theory write the HJB equation in the form

$$\rho V = u(c) + \mathcal{A}V,$$

where $\mathcal{A}$ is a special differential operator associated with the control problem called the *infinitesimal generator*. The operator $\mathcal{A}$ is the continuous time analog of the transition matrix for a discrete system, mapping a value function to its continuation value. The matrices $\mathbf{A}_m$ approximate $\mathcal{A}$. There is a deep connection between the HJB and KF equations through $\mathcal{A}$ as discussed earlier.

the next time step is tighter, this may not be enough: by saving zero, the agent may suddenly find himself in violation of the borrowing constraint. Instead, the borrowing constraint needs to force the agent to save a strictly positive amount, enough to take him from the current borrowing constraint to that of next period.

Hence, if $\underline{w}^n$ denotes today's wealth constraint and $\underline{w}^{n+1} > \underline{w}^n$ denotes the next one, a special boundary condition needs to be applied at each grid point between $\underline{w}^n$ and $\underline{w}^{n+1}$ (inclusive). The boundary condition ensures that savings are high enough for next period's expected wealth to equal $\underline{w}^{n+1}$.

Upper wealth constraint

Implicit in the idea of solving this system by discretizing it onto a grid is the requirement to bound the state variables. In particular, this means we need an upper limit $\bar{w}$ on wealth. In order to do so in a model-consistent way, we impose $\bar{w}$ as a reflecting boundary on the wealth process w_t, just like the lower boundary, which forces savings to be nonpositive at that point. Note that this boundary condition does not bind and will not be selected by the upwind method, and so should not affect the model in an economically significant way.

Finally, a second-order boundary condition needs to be imposed at $\bar{w}$, in order to pin down capital holdings at those points. (Note that we did not need to do this at $\underline{w}$ because the conditions $k \geq 0$ and $k \leq w - \underline{w}$ ensure that $k = 0$ at the lower wealth boundary.) We follow Achdou et al. (2017) in imposing

$$\partial_{ww} V(t, \bar{w}, z) = -\gamma \frac{\partial_w V(t, \bar{w}, z)}{\bar{w}}. \tag{1.47}$$

This condition results follows from the asymptotic behavior of the value function for large w, as derived in Lemma 1.1.

KF equation

As outlined by Achdou et al. (2017), we solve equation (KF) using a finite difference method. It turns out that if one chooses the approximations to the derivatives judiciously, we can exploit the adjointness of the HJB and KF equations and solve the KF equation for free (see Footnote 27 – for more details, please see Achdou et al. (2017, online appendix)):

$$\frac{g^{n+1} - g^n}{\Delta t} = (\mathbf{A}_n^T - \lambda_n \mathbf{I}) g^{n+1}. \tag{1.48}$$

Here $\mathbf{A}_n^T$ is the transpose of the transition matrix obtained on the convergence of the Newton method for the HJB equation at time point t_n, and $\lambda_n I$ takes into account the death rate at t_n.

This implicit method pins down g^{n+1} from g^n, and we can solve it efficiently using sparse matrix methods.

This defines a forward-looking solution scheme for the density of wealth at each age. It is solved starting from a density g^0 which is the initial distribution of wealth (i.e. inheritance).

Stationary equilibrium

To sum up, the algorithm to solve for a stationary equilibrium is as follows.

1. Make a guess of the stationary interest rate r
2. Solve the HJB equation backwards given r using the fully implicit method to obtain agents policies
3. Solve the KF equation forward to obtain the wealth distribution

4. Using the two steps above, work out the aggregate excess demand for bonds given r

5. Increase the guess of r if there is excess supply and decrease it if there is excess demand

6. Repeat until the bond market clears.

1.B.2 *Transition equilibrium*

The algorithm to solve for a transition equilibrium in this model requires an extra layer compared to transition in the Achdou et al. (2017) model: every point in time contains agents of all ages. Due to lifecycle effects, at time τ we need to solve the agent's problem, a time-dependent system, for each cohort with date of birth $s \leq \tau$. By solving for the decision functions at each age for each cohort, we solve the corresponding KF equations to get the density functions of wealth at each age for each cohort. By selection the density functions for agents of cohort s with age $\tau - s$, we can piece together the wealth distribution at each time τ.

When agents learn of an unexpected shock, they change their policy functions from the initial steady state. This is captured by a change in the transition matrix $\mathbf{A}(s, \tau)$. This way, the KF equations change, capturing the shock.

Transition algorithm

Denote transition calendar time by τ. We consider transition experiments from time $\tau = 0$ to some maximal time $\tau = \bar{\tau}$.

For each $\tau \in [0, \bar{\tau}]$ let

$$X_\tau = (D_\tau, \rho_\tau, \sigma_\tau, \alpha_\tau, vol_\tau, \ldots, (\phi^s_\tau)_{s \leq \tau}, (\lambda^s_\tau)_{s \leq \tau}, (\chi^s_\tau)_{s \leq \tau}, (y^s_\tau)_{s \leq \tau}, (BC^s_\tau)_{s \leq \tau})$$

be the tuple consisting of all model parameters (demographic and economic) at time τ, for every cohort born before time τ. We assume stationarity before 0 and after $\bar{\tau}$, so that $X_{-\tau} = X_0$ and $X_{\bar{\tau}+\tau} = X_{\bar{\tau}}$ for all $\tau > 0$.

The following sketches the algorithm used to solve for transition equilibria.

1. Discretize the transition time interval $[0, \bar{\tau}]$ onto a grid $\boldsymbol{\tau} = \tau_0, \ldots, \tau_N$.

2. Guess a time path for the interest rates $(r_\tau)_{\tau \in \boldsymbol{\tau}}$

3. At any time $\tau \in \boldsymbol{\tau}$ there will be agents of every age. We need to know how they make decisions. Hence, at each time $\tau \in \boldsymbol{\tau}$, we need to solve the HJB equation for *each cohort* $s \leq \tau$. Assuming an upper bound on lifetimes T, this means that we need to solve it for finitely many cohorts: each date of birth $s \in [\tau - T, \tau] \cap \boldsymbol{\tau}$.

   For each cohort s, work out the age-dependent paths (i.e. lifecycle profiles) of each parameter and endogenous variable for every cohort of agents born at s. These are encoded into $(\tilde{X}^s_t)_{t \geq 0} = (X_{s+t})_{t \geq 0}$. Use the lifecycle paths $(\tilde{X}^s_t)_{t \geq 0}$ for cohort s to solve the HJB equation for each cohort. These are solved using the fully implicit Newton method to solve the lifecycle problem as described earlier.

4. The next step is to calculate wealth distributions, but we need to be careful when doing this as agents' policy functions change when the shock is announced. At first, each cohort's age-dependent wealth distribution will evolve using the transition matrix for the initial steady state. When the shock occurs, agents' decisions change, so the densities now evolve with a new transition matrix.

Recall the numerical scheme for obtaining the age- and cohort-specific wealth distributions outlined in Section 1.B.1:

$$\frac{g^{n+1} - g^n}{\Delta t} = (A_n^T - \lambda_n I)g^{n+1},$$

where $\mathbf{A}_n^T$ is the transpose of the transition matrix for the value function obtained in solving the HJB equation with the Newton method. Update the age-dependent wealth distribution using the appropriate transition matrix for each cohort.

5. With the age-dependent wealth distribution for each cohort $g(-; s, t)$, we can use the age structure $D(t, \tau)$ to obtain the unconditional distribution as usual.

   At each time step τ, we use the policy functions for each cohort together with the age-dependent wealth distribution to work out aggregate wealth and capital holdings for each cohort. By aggregating up using the demographic distribution, we obtain unconditional aggregate wealth W_τ and capital holdings K_τ. Excess demand for bonds at each point in the transition is given by

   $$S_\tau = W_\tau - K_\tau.$$

6. Given the excess demand for bonds S_τ at each time, update the guess of $(r_\tau)_\tau$ to

   $$r - \xi \cdot S$$

   for some vector of weights ξ with positive entries, a form of Walrasian tâtonnement. These weights can be determined through trial and error, or endogenously using a quasi-Newton method (e.g. the secant method).

7. Go back to step 2 and repeat until $(r_\tau)_\tau$ converges and $S_\tau \approx 0$ for all t.

Although this algorithm contains 4 large nested loops, at the inside of which one solves a big matrix inversion problem, the process usually converges in around 20 minutes on a good computer for reasonable grid sizes.

1.B.3 *Stability and convergence of numerical scheme*

As discussed in Section 1.A.2, the theory has not fully caught up with techniques in this field yet: although there exists a proof of existence of a transition solution of the system, uniqueness results are still being sought.

When it comes to the numerical methods, I rely on Barles and Souganidis (1991). They show that the finite difference schemes of the kind we employ to solve the HJB equation are guaranteed to converge to the unique (viscosity) solutions of these equations under mild conditions.

Part II

CHAPTER 2

FAT TAILS IN WEALTH FROM FAT TAILS IN AGE?

2.1 INTRODUCTION

Top wealth inequality has increased dramatically in the United States since the 1970s. Figure 2.1 shows the evolution of the *tail indexes* of the wealth distribution, a measure inversely related to top wealth inequality, at two points in the wealth tail (between the top 1% and 10% and the top 0.1% and 1%). This plot suggests two stylized facts about changes in the U.S. wealth distribution over the past decades.[1]

Firstly, the wealth distribution had a relatively thin, Gaussian, tail through

[1]The empirical tail index at the 1%-10% level is defined as

$$\hat{\zeta}_{1\%-10\%} = \frac{1}{\log_{10}\left(\frac{x_{1\%}}{x_{10\%}}\right) + 1},$$

where $x_{1\%}$ is the share of wealth held by the 1% top quantile. If there is little tail inequality and $x_{1\%}$ is a small fraction of $x_{10\%}$, the tail index is larger.

The tail index is linked to the Pareto distribution, a fat tailed distribution which has a density function of the form

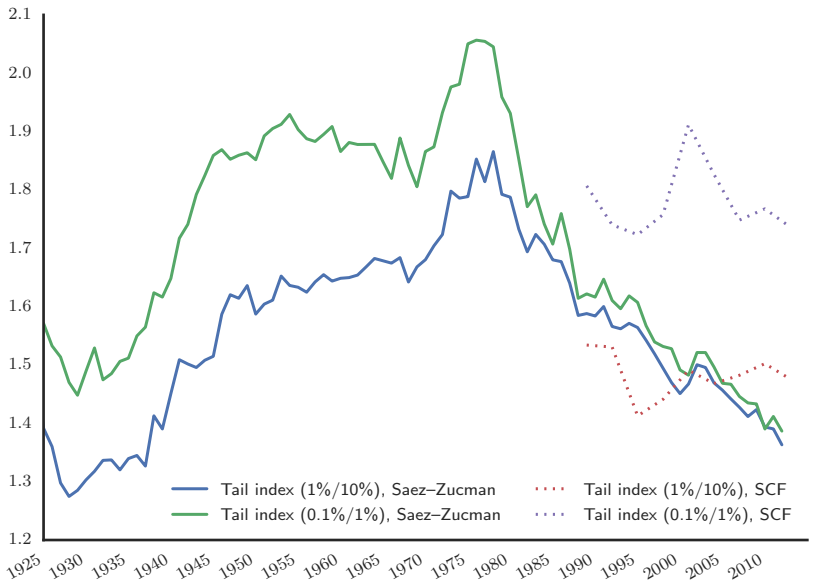
$$f(x) = kx^{-\zeta}.$$

The parameter $\zeta > 1$ captures the degree of inequality in the distribution, where lower ζ corresponds to higher inequality. This distribution has been shown closely to fit the distribution of many random growth processes such as firm size, city size, and even the number of peas in a pod. It is named after Vilfredo Pareto, who first used it in modeling the income distribution in Italy in 1896 (Pareto 1896). Gabaix (2009) provides a great overview of how and where these distributions appear in economics and finance.

The tail of the Pareto distribution is characterized by a fractal property: if the top 0.01% of the population is ϵ times richer than the top 0.1%, then the top 0.1% must in turn be the same ϵ times richer than the top 1%, and so on. As a result, for a Pareto distribution, $\zeta = \hat{\zeta}_{1\%-10\%} = \hat{\zeta}_{0.1\%-1\%}$. The value of ϵ depends on ζ, providing a criterion we can take to the data. This fractal property is strongly violated in a thinner-tailed distribution such as the log-normal, in which log tail density declines quadratically.

FIGURE 2.1: TAIL WEALTH INEQUALITY IN THE US

Note: Figure plots estimates of the Pareto tail index ζ for wealth over time using percentile ratios, using data from Saez and Zucman (2016) and from the Survey of Consumer Finances. Lower ζ corresponds to higher tail inequality. In a Pareto (fat tailed) distribution, these lines would coincide.



most of the 20th century, but after 1980 its shape changed into a fat tail: until 1980, tail inequality was lower as one looked deeper into the tail, but after the 1980s this relationship ceased to hold. The fact that the green and blue lines seem to coincide after the 1980s is suggestive of convergence to a Pareto distribution for the wealth tail.[2] The Pareto distribution has a fatter tail than Gaussian distributions, and arises in many areas of the natural sciences, economics, and finance.

Secondly, in the decades since 1980, the wealth tail, already Pareto in shape,

[2]Note that the estimates from SCF data seem to tell another story. See Kopczuk (2015) and Bricker et al. (2016) for discussion of this discrepancy.

has gotten fatter. Wealth inequality in the United States has increased quickly: estimates of the tail index have fallen by 20% over that period.

Benhabib et al. (2011, 2015, 2016) argue that the fat Pareto wealth tail observed today can be explained by idiosyncratic risk in household returns to capital. The idea is that around half of household wealth in the U.S. is invested in private residence equity or in private business equity, and each of these carries high amounts of idiosyncratic risk. As a result, the wealth tail is constituted by agents who have been fortunate enough to have long periods of high returns.

In this chapter, I argue that *although this mechanism does deliver a fat Pareto tail in the long run, its dynamics are slow.* I prove that in a context with realistic demographics, it operates on the order of generations, depending crucially on bequests and inheritance. This means that this mechanism would explain the prevailing fat tail in the wealth distribution using centuries-old dynasties. However, given that the wealth tail in the 1970s was much thinner, these ancient dynasties would have had to appear suddenly in the 1980s.

The results derived in Benhabib et al. (2016) emphasize this point further. The authors show that a perpetual youth model yields a steady-state wealth distribution with a fat Pareto tail, even if the only source of uncertainty is lifespan. There, households accumulate wealth deterministically as they age. The fat wealth tail consists exclusively of households that are centuries old.

When the mechanism is instead studied in a model derived from Chapter 1, featuring realistic lifespans, these agents are eliminated and the wealth tail becomes much thinner.

It is also difficult for the slow transition dynamics of idiosyncratic returns to capital to explain the second stylized fact, in which the tail index of the

wealth distribution has increased since the 1980s. This point was made by Gabaix et al. (2016) using a stylized model. I confirm their results with a numerical experiment using the larger general equilibrium framework of Chapter 1.

Gabaix et al. (2016) provides a study of the dynamics of inequality using different approaches. Idiosyncratic returns to capital lead household wealth to be a random growth process, and they prove that such processes have slow speed of convergence.

The idea that this mechanism would explain current levels of wealth inequality can be thought of as analogous to Luttmer (2011). Focusing on the fat-tailed distribution of firm size, he argues that commonly proposed random-growth processes for firm size would counterfactually imply centuries-old firms.

Benhabib et al. (2019) provide an analysis of several different proposed mechanisms for explaining the Pareto wealth tail, and find them all – including idiosyncratic capital returns – to be quantitatively important. This underlines the idea that although a different mechanism should be used to explain the dynamics,[3] idiosyncratic capital returns are still useful for understanding long-run properties of the wealth distribution.

2.2 DEMOGRAPHICS AND IDIOSYNCRATIC CAPITAL RETURN RISK

As discussed above, we can draw two conclusions from the data about the wealth distribution in the United States since 1980 by looking at Figure 2.1. Firstly, that the wealth tail has transition from a thin tail to a fat Pareto tail, and secondly that inequality has increased rapidly.

[3]Gabaix et al. (2016) argue that a mechanism that departs from random growth assumptions would be needed, such as superstar entrepreneur theories.

In this section, I discuss how well the idiosyncratic returns to capital (Benhabib et al. 2011) can explain these phenomena. I will show that the dynamics of this mechanism are slow, relying on intergenerational wealth transfers.

Benhabib et al. (2011, 2015, 2016) show that with this risk in capital returns, the stationary distribution of wealth has a Pareto fat tail, even if the initial wealth distribution is thin-tailed. In the first two studies, households are either infinitely lived, or equivalently are members of a dynasty dating back to $-\infty$, inheriting the wealth of their ancestors. These ideas seem at odds with Figure 2.1, which shows that the wealth distribution goes through periods of high and low inequality and rapid transition from thin to fat tail. Indeed, historically, very few families have avoided the various wars and redistributive regimes of the past few centuries (not to mention irresponsible inheritors) with their assets unscathed. If convergence of the wealth distribution to a Pareto is slow in these models, not enough time has elapsed since the period of low inequality observed in the 1970s until today for this mechanism to explain the seeming convergence of the empirical wealth tail to a Pareto.

Benhabib et al. (2016), by contrast, do obtain a stationary distribution of wealth with a Pareto tail in a model where agents have finite lifespans, even absent intergenerational transfers. In this perpetual youth model, there is a fat right tail in the distribution of *ages* in steady state, so that there is a large mass of extremely old agents present in the economy. It turns out that it is these agents, who have been compounding interest for often hundreds of years, who drive the fat wealth tail in this model. Indeed, the authors show that the wealth distribution in a perpetual youth model retains a Pareto tail even when all sources of household risk are disabled.

To understand this mechanism and how it changes in a demographic model, let us introduce a simple model of household wealth.

Let $x_t^s = \log w_t^s$ denote the logarithm of wealth at time t for an agent born at time s. Suppose that $(x_t^s)_t$ follows a Brownian motion with drift for $t \geq s$, and that $x_s^s = 0$, so that

$$\mathrm{d}x_t^s = \mu \,\mathrm{d}t + \sigma \,\mathrm{d}B_t^s, \tag{2.1}$$

where B_t^s are independent (over s) Brownian motions, which start at $B_s^s = 0$.

This setup captures the features relevant to our analysis of the full model with idiosyncratic capital returns of Chapter 1. We saw in Section 1.A.1 that in that model, agents' savings policy functions are approximately linear in wealth when wealth is high. This implies that in the tail, the wealth process evolves like a geometric Brownian motion, or equivalently that log-wealth follows a Brownian motion in the tail. This corresponds to assuming equation (2.1) for suitably chosen parameters μ and σ. Although the policy functions are only linear at high levels of w, i.e. strictly speaking, the dynamics of x only behave like a Brownian motion for large x, I assume in this section that these dynamics hold for all x. The reason for this is that we are only interested in modeling the tail distribution of x, and so we do not lose anything by making this simplifying assumption.

The main intuition for these derivations is that in a finite amount of time, a Brownian motion can only diffuse its initial distribution by so much. If it started out thin-tailed (and starting it at a fixed point is an extreme example of this), it will take it infinitely long to diffuse enough to generate a fat tail.

As a Brownian motion, log-wealth x_t^s follows a Gaussian distribution with

mean $(t-s)\mu$ and variance $(t-s)\sigma^2$. Thus,

$$f(x;s,t) = \frac{1}{\sqrt{2\pi\sigma^2(t-s)}} \exp\left(-\frac{1}{2(t-s)\sigma^2}(x-(t-s)\mu)^2\right) \tag{2.2}$$

is the density function for log-wealth among agents of cohort s at time t.

Let $D(s,t)$ denote the age structure, the density function for agents of age s alive at time t. Integrating $f(x;s,t)$ against $D(s,t)$ yields the unconditional distribution of log-wealth at time t, function $g(x;t)$,

$$\begin{aligned} g(x;t) &= \int_{-\infty}^{t} f(x;s,t)\hat{D}(s,t)\,\mathrm{d}s \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^{\infty} \frac{1}{\sqrt{\tau}} \exp\left(-\frac{1}{2\sigma^2\tau}(x-\mu\tau)^2\right) D(\tau,t)\,\mathrm{d}\tau, \end{aligned} \tag{2.3}$$

where in the second line re-indexes the integral from cohort to age $\tau = s-t$ (and uses an accordingly re-indexed demographic function $D(\tau,t) = \hat{D}(t-\tau,t)$).

We are interested in the properties of function $g(x;t)$ for large x. Note that the age-specific distribution of log-wealth $f(x;s,t)$ has a Gaussian tail. As age $s-t$ increases, so does the variance of x_t^s.

As we shall see, in a perpetual youth model, $D(\tau,t)$ places large mass on old agents. As a result, the weighted average equation (2.3) emphasizes these agents with high variance. This ends up yielding a Pareto tail.

All proofs can be found in Section 2.A. It is useful for the remainder of the section to define terms.

Definition 2.1. *Call a density function $h(y)$* fat-tailed *in y if there exists $A > 0$ such that*

$$\lim_{y\to\infty} \frac{h(y)}{e^{-Ay}} > 0.$$

In other words, $h(y)$ declines to zero no slower than an exponential decay as y grows.

For example, the Pareto distribution is fat-tailed, but not the Gaussian distribution.

The following result demonstrates that if there is an upper limit to agents' lifespans, the distribution of wealth has a thin tail. In other words, while a perpetual youth model can deliver a Pareto tail even if all agents start with the same wealth, this cannot be the case when agents' lifespans are bounded.

PROPOSITION 2.1. *If $D(\tau, t)$ has bounded support (in age τ), and initial wealth is fixed $x_s^s = 0$, then the unconditional distribution $g(x;t)$ is dominated by a Gaussian function in the tail at all times t.*

In fact, the following proposition establishes a stronger result. It turns out that perpetual youth is precisely in the knife-edge region of thinnest-tailed demographic structures that still deliver a fat-tailed wealth distribution.

PROPOSITION 2.2. *Suppose all agents have $x_0 = 0$. Then for all t, the unconditional distribution of log-wealth $g(x;t)$ is fat tailed (in x) if the demographic distribution $D(\tau, t)$ is fat-tailed (in τ).*

As a partial converse, if D is asymptotically dominated by a Gaussian in τ, then so is g by a Gaussian in x.

Hence, even when the ages are not bounded, as long as the density of ages declines at Gaussian speed, the wealth tail will be thin tailed if initial wealth is deterministic. Conversely, we can apply the theorem to prove the result of Benhabib et al. (2016).

COROLLARY 2.1. *In a model with perpetual youth demographics, idiosyncratic capital returns lead to a stationary Pareto distribution of wealth.*

This corollary follows from the fact that the age structure in a perpetual youth model with replacement is exponential,

$$D(\tau) = \lambda e^{-\lambda\tau},$$

where λ is the mortality rate.

Asymptotics

Due to the probabilistic nature of agents' lifespans, there is a small but positive probability of agents reaching any given age. We can ask what the distribution of log-wealth converges to as age tends to infinity, conditional on survival.

It turns out that the distribution wealth within a cohort converges to a Pareto distribution in the limit as age $\tau \to \infty$, conditional on survival.[4] This can be seen by guessing and verifying a Pareto density function into the Kolmogorov Forward equation for a cohort, equation (KF).

In fact, this result is not surprising. Gabaix (2009) explain how a random growth process with some kind of *stabilizing force* will have a stationary distribution with a Pareto tail. Here, wealth is a random growth process (its growth rate is determined by the random returns to capital), and it is constrained from below by a reflecting boundary at the borrowing constraint.

This mechanism underlies the results of Benhabib et al. (2011, 2015, 2016). The unconditional distribution is a weighted average of the cohort conditional densities, which are increasingly fat in the tail. Unless the demographic structure

[4]This statement assumes that the agent's problem becomes stationary as age increases. In other words, at a sufficiently advanced age, the agent's problem collapses to a perpetual youth type problem. This is indeed the assumption made for advanced old age in Chapter 1.

decays sufficiently fast with age, there will be undue weight on these large-tailed wealth distributions.

It is worth noting that Benhabib et al. (2011) uses a discrete time approach, deriving properties stationary wealth distribution with use of the stochastic difference equations, Kesten processes, and the results of Saporta (2004, 2005). In a continuous time setting, these proofs can be shortened considerably by use of the Kolmogorov Forward equation.

Inheritance

Until now, we have assumed that the initial wealth of entering households is fixed in order to derive our results. More generally, the result holds if initial wealth has any thin-tailed distribution:

PROPOSITION 2.3. *Suppose agents begin life with log-wealth x_0, drawn from a density function $f_0(x)$. Suppose further that the distribution of ages $D(\tau, t)$ in the economy has bounded support in τ. Then at any time t, the unconditional wealth distribution,*

$$g(x;t) = \int_0^\infty f(x;\tau,t)D(\tau,t)\,\mathrm{d}a,$$

is fat tailed if and only if $f_0(x)$ is.

We can use this proposition to think about dynasties. Suppose the first generation of a dynasty began life with wealth drawn from a thin tailed distribution. Throughout their lifetime, the wealth of agents from this first generation would evolve like a geometric Brownian motion. As the policy functions are linear, their choice of bequest would also follow a geometric Brownian motion. It follows therefore that the distribution of wealth inherited by their successors would

be thin tailed as well, although with a higher variance than their ancestors, due to the accumulated variance in returns on wealth over the agent's lifetime.

By this inductive line of reasoning we conclude that unless there is an unbroken chain of bequest and inheritance dating back to $-\infty$, such an economy would display a wealth tail that is not Pareto.

However, although the distribution never actually reaches a Pareto tail, it *is* converging to one. It may be that in finite time the wealth tail is all but indistinguishable from a Pareto. For this reason, I run numerical experiments in the next section in order to quantify the wealth tail and see just how far it is from a Pareto.

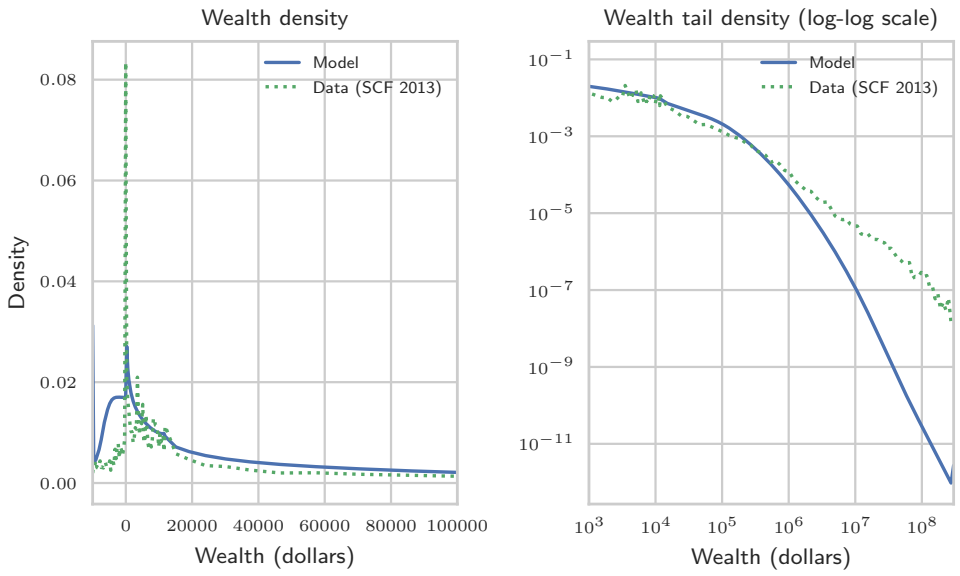
2.3 quantifying the tail of the wealth distribution

In the previous section I argued that with idiosyncratic risk to capital returns, convergence of the wealth distribution to a fat tail is slow, relying on intergenerational transmission. If the distribution of initial wealth is thin-tailed – as was the case in the 1970s – then the overall wealth tail cannot be Pareto. However, the possibility remains that although mathematically thin-tailed, the distribution is numerically difficult to distinguish from a Pareto fat tail. In this section, I demonstrate by means of a calibrated model that this is not the case: the distribution of wealth is counterfactually thin, with top wealth shares too low relative to the data.

I use the calibrated overlapping generations model presented in Section 1.3. In this model, there is no intergenerational transmission of wealth. Instead, the distribution of wealth of new households is thin-tailed (lognormal), targeting

FIGURE 2.2: DISTRIBUTION OF WEALTH IN THE MODEL AND THE DATA

Note: Figure plots wealth density in levels (left) together with its tail in log-log scale (right), for the model and the 2013 Survey of Consumer Finances



two empirical moments for 25 year-old households in the Survey of Consumer Finances: the mean and variance of log wealth.

The idea is to start the model with a thin-tailed distribution (as was prevalent in the 1970s), and to see after one generation how fat is the resulting distribution, comparing it to the tail found today. We shall see that in this case the model produces insufficient tail wealth, suggesting that idiosyncratic returns to capital are not sufficient for explaining the changes in the U.S. wealth distribution since the 1970s.

Figure 2.2 plots the density function for wealth in the model, against that calculated from the Survey of Consumer Finances for the United States in 2013. The model matches several empirical features, including the spike around zero, the skewness and the long right tail.

Table 2.1: Steady state wealth distribution statistics.

	Model	US SCF 2013	US SCF 1989
Top 10% wealth share	58%	75%	67%
Top 1% wealth share	21%	36%	30%
Top 0.1% wealth share	6.2%	13.5%	10.8%
Implied tail index (1%-10%)	1.53	1.48	1.53
Implied tail index (0.1%-1%)	2.11	1.74	1.80

The second panel in the figure plots the tail of the wealth distribution in log-log scale. Since a Pareto density function is proportional to $x^{-\zeta-1}$, a Pareto tail would manifest as a straight line in this plot. A thin tail would be concave, with a Gaussian appearing as a negative quadratic. As explained in the last section, we see that the model tail is indeed thin, and looks Gaussian. By contrast, tail density in the data looks much more Pareto.

The fit of the model wealth tail seems to drop off after the million dollar mark (corresponding to roughly the top decile of the empirical distribution).

Table 2.1 lists further statistics on the tail wealth distribution in the model and compares them to the data. The model undershoots the empirical top wealth shares.

The bottom rows are two estimates of the tail index ζ if the distribution were Pareto, calculated using the relation

$$1/\zeta = \log_{10}(x_{1\%}/x_{10\%}) = \log_{10}(x_{0.1\%}/x_{1\%})$$

linking it to top percentile levels.[5] This equality is close to holding in the data, suggesting a Pareto empirical tail. However, this is not the case for model esti-

[5]This equation derives from the distribution function for a Pareto distribution, $1 - F(x) \propto x^{-\zeta}$, applied to the top percentiles $x_{1\%}$ etc.

mates: although the 1%–10% estimate is good, the 0.1%–1% number is too high, corresponding to a thin tail with insufficient top tail inequality in the model.

To sum up, we see that starting from a thin-tailed distribution of wealth as was prevalent in the 1970s, the resulting wealth distribution one generation later remains badly thin-tailed relative to the data. A model with idiosyncratic risk to capital returns could only match current levels of top wealth inequality if initial wealth were already quite fat in the tail.

2.4 CONCLUSION

This chapter examines a commonly proposed mechanism for explaining the fat tail of the U.S. wealth distribution using the framework developed in Chapter 1. I argue that this mechanism, idiosyncratic risk to household investment returns, has dynamic properties that put it at odds with changes observed in the U.S. wealth distribution over the past few decades. Using a mix of theoretical and numerical results, I show that this mechanism is slow, with its dynamics operating on the order of generations. As a result, this channel alone cannot account for both the relatively thin-tailed distribution observed in the 1970s and the fat tail seen after 1980.

Further, using a calibrated model featuring idiosyncratic risk to capital returns and overlapping generations, I show that this mechanism does not noticeably thicken the wealth tail after one generation. Measures of top wealth inequality in the model badly undershoot the data.

These observations highlight the importance of modeling lifecycle savings behavior and intergenerational wealth transmission in understanding the distribu-

tion of wealth.

Gabaix et al. (2016) argue that mechanisms such as this which lead to a random rate of growth in household wealth have slow convergence speeds and will struggle to match the dynamics observed in the data. This suggests that another mechanism which departs from this assumption, such as theories of superstar investors, may help us understand changes in the U.S. wealth distribution.

In future research, I will evaluate these alternative mechanisms numerically in a similar model, in order to examine their ability to explain the dynamics of the U.S. wealth distribution. A candidate theory that would fit these requirements is that changes in the tax and transfer system in the United States (documented by Piketty and Saez (2007)) may have led to persistent heterogeneity in returns. Although studies have examined the extent to which this may explain the level of top wealth inequality, a study of the dynamics of wealth in this setting would be welcome.

2.A PROOFS

Recall the setup of Section 2.2: if the log-wealth process x_τ of each agent is started at zero, then the density of wealth of agents of age τ is Gaussian,

$$\begin{aligned} f(x;\tau) &= \frac{1}{\sqrt{2\pi\sigma^2 a}} \exp\left(-\frac{1}{2\sigma^2 a}(x-\mu\tau)^2\right) \\ &= M(x)\frac{1}{\sqrt{\tau}} \exp\left(-\frac{x^2}{2\sigma^2\tau}\right) \exp\left(-\frac{\mu^2\tau}{2\sigma^2}\right) \end{aligned} \tag{2.4}$$

where

$$M(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{x\mu}{\sigma^2}\right) \tag{2.5}$$

depends only on x.

The unconditional density of wealth $g(x;t)$ is obtained by integrating the conditional distributions $f(x;\tau)$ against the demographic distribution:

$$g(x;t) = \int_0^\infty f(x;\tau)D(\tau,t)\,\mathrm{d}\tau.$$

Proof of Proposition 2.1. Suppose that $D(\tau,t)$ has bounded support in τ, with $T = \sup\{\tau > 0 : D(\tau,t) > 0 \text{ for any } t > 0\}$ being the upper bound. Then

$$\exp\left(-\frac{x^2}{2\sigma^2\tau}\right) \leq \exp\left(-\frac{x^2}{2\sigma^2 T}\right)$$

D-almost surely, so

$$\begin{aligned} g(x;t) &= \int_0^\infty f(x,\tau)D(\tau;t)\,\mathrm{d}\tau \\ &\leq M(x)\exp\left(-\frac{x^2}{2\sigma^2 T}\right)\int_0^\infty \frac{1}{\sqrt{\tau}}\exp\left(-\frac{\mu^2\tau}{2\sigma^2}\right)D(\tau;t)\,\mathrm{d}\tau. \end{aligned}$$

It is easy to see that the final integral converges. The Gaussian term $\exp\left(-\frac{x^2}{2\sigma^2 T}\right)$ asymptotically dominates $M(x)$, so that $g(x;t)$ itself is asymptotically dominated by a Gaussian. □

Before proving Proposition 2.2, we need to establish a few lemmas.

LEMMA 2.1. *Suppose that $D(\tau,t)$ is Gaussian in τ. Then*

$$g(x;t) = \int_0^\infty f(x;\tau)D(\tau,t)\,\mathrm{d}\tau$$

is dominated by a Gaussian.

Proof. Let $\bar{D}(\tau,t) = \exp\left(-\frac{\mu^2\tau}{2\sigma^2}\right)D(\tau,t)$. Then $\bar{D}$ is also Gaussian, and we can write

$$g(x;t) = M(x)\int_0^\infty \underbrace{\frac{1}{\sqrt{\tau}}\exp\left(-\frac{x^2}{2\sigma^2\tau}\right)}_{P(x;\tau)}\bar{D}(\tau,t)\,\mathrm{d}t. \tag{2.6}$$

For all x, there is a value $T(x)$ such that $\tau \mapsto P(x;\tau)$ is concave for all $\tau > T(x)$. We can write

$$\bar{D} = \bar{D}1_{[0,T(x)])} + \bar{D}1_{[T(x),\infty)},$$

and split up the integral in equation (2.6) accordingly. One of the terms is an integral against a distribution with bounded support, for which we can apply Proposition 2.1.

The second term has an integral of a strictly concave function against a (truncated) Gaussian. Jensen's inequality then bounds this second term by a Gaussian, and we are done. □

We have thus seen that a Gaussian age structure leads to a wealth distribution with a Gaussian tail. In the following lemma we see a partial converse to this statement: a Pareto *age* distribution implies a wealth distribution with a Pareto tail. (Recall that a Pareto distribution for wealth is equivalent to an Exponential distribution for log-wealth.)

Lemma 2.2. *Suppose* $D(\tau,t) = \lambda\exp(-\lambda\tau)$ *for some* $\lambda > 0$. *Then*

$$g(x;t) = \int_0^\infty f(x;\tau)D(\tau,t)\,\mathrm{d}\tau$$

is asymptotically proportional to $\exp(-Ax)$ *for some* $A > 0$.

Proof. We can write $g(x;t)$ as

$$g(x;t) = M(x)\lambda\int_0^\infty \frac{1}{\sqrt{\tau}}\exp\left(-\frac{x^2}{2\sigma^2\tau}\right)\exp\left(-\frac{\mu^2\tau}{2\sigma^2}\right)\exp(-\lambda\tau)\,\mathrm{d}\tau. \qquad (2.7)$$

We simplify this integral using the symmetry of the Laplace transform: for two functions $u, v \in L^2$,

$$\int uv = \int (\mathcal{L}u)(\mathcal{L}^{-1}v).$$

Write the integrand in equation (2.7) as

$$\underbrace{\frac{1}{\sqrt{\tau}}\exp\left(-\frac{x^2}{2\sigma^2\tau}\right)}_{P(\tau;x)}\underbrace{\exp\left(-\left(\frac{\mu^2}{2\sigma^2}+\lambda\right)\tau\right)}_{Q(\tau)}.$$

It turns out that $P(\tau; x)$ has an analytical Laplace transform,

$$(\mathcal{L}P)(s) = \frac{\sqrt{\pi}}{\sqrt{s}}\exp\left(-\frac{x}{\sigma}\sqrt{s}\right).$$

Function $Q(\tau)$ is the Laplace transform of a delayed impulse, and so we can look up its inverse transform

$$(\mathcal{L}^{-1}Q)(s) = \delta\left(s-\left(\frac{\mu^2}{2\sigma^2}+\lambda\right)\right),$$

where $\delta(-)$ is the Dirac delta.

Putting these together, we see that

$$\begin{aligned} g(x;t) &= M(x)\lambda\int_0^\infty (\mathcal{L}P)(s)(\mathcal{L}^{-1}Q)(s)\,\mathrm{d}s \\ &= M(x)\lambda\left[(\mathcal{L}P)\left(\frac{\mu^2}{2\sigma^2}+\lambda\right)\right] \\ &= M(x)\lambda\frac{\sqrt{\pi}}{\sqrt{\frac{\mu^2}{2\sigma^2}+\lambda}}\exp\left(-\frac{x}{\sqrt{\sigma^2\tau}}\sqrt{\frac{\mu^2}{2\sigma^2}+\lambda}\right). \end{aligned} \tag{2.8}$$

The second equality follows from the fact that $\mathcal{L}^{-1}Q$ is a Dirac delta, and the third from the definition of $\mathcal{L}P$.

The function $M(x)$ is $O(e^x)$, as is the term on the right-hand side of equation (2.8), so we are done unless they exactly cancel (which will only happen for one knife-edge value of λ, which we control). □

Finally, we are ready to prove the main theorem.

Proof of Proposition 2.2. Lemma 2.1 implies that if $D(\tau, t)$ is dominated by a Gaussian then the wealth distribution

$$g(x;t) = \int_0^\infty f(x;\tau)D(\tau,t)\,\mathrm{d}\tau.$$

is thin-tailed. Similarly, Lemma 2.2 implies that if D has tails as fat (or fatter) than an exponential distribution, then so does g.

It remains to deal with the intermediate cases. □

Proof of Proposition 2.3. Suppose x_τ denotes the process for log-wealth started deterministically at zero, and $\hat{x}_\tau$ denotes the process for log-wealth started from some density function $f_0(x)$. Then

$$\hat{x}_\tau = x_\tau + \hat{x}_0,$$

and the two terms on the right hand side are independent, due to the independent increments of Brownian motion. Then the density $\hat{f}$ of $\hat{x}_\tau$ is the convolution

$$\hat{f}(x,\tau) = \int_0^\infty f(x - x_0, \tau)f_0(x_0)\,\mathrm{d}x_0 \qquad \text{for all } \tau > 0.$$

Claim 2.1. *If f is a Gaussian function, then the convolution*

$$\hat{f}(x,\tau) = \int_0^\infty f(x - x_0, \tau)f_0(x_0)\,\mathrm{d}x_0 \qquad \textit{for all } \tau > 0,$$

is dominated by a Gaussian in the tail if and only if the function $\hat{f}_0$ is dominated by a Gaussian.

This can be seen by taking Fourier transforms, which maps convolutions to multiplications, and by noting that Fourier transforms take Gaussians to Gaussians.

With this claim in hand, we know that $\hat{f}$ is dominated by a Gaussian. We can then apply the proof of Proposition 2.1 and we are done. □

CHAPTER 3

INSPECTING THE MECHANISM

3.1 INTRODUCTION

This chapter consists of three free-standing sections, each exploring additional applications and mechanisms of the overlapping generations heterogeneous agent framework developed in Chapter 1. I show that lifecycle motives alter household savings and portfolio decisions, revealing feedback mechanisms between distributional changes and aggregates through asset price changes.

In Section 3.2, I show that lifecycle motives account for a large fraction of aggregate savings in this model, with implications for asset prices and aggregates. This exacerbates the overaccumulation of capital attributed to precautionary saving against income fluctuations, first identified by Huggett (1993) and Aiyagari (1994).

Section 3.3 analyzes different sources of secular stagnation in this framework. I quantify the macroeconomic and distributional effects of lengthening lifespans and rising structural income inequality, and analyze implications for fiscal policy. Both trends raise aggregate household savings and lower safe interest rates, but the distributional effects are mixed.

The decline in the equilibrium interest rate itself has distributional implications, which I study in Section 3.4, a partial converse to Section 3.3. In this section, I examine the distributional impact of a permanent change in the risk-

free rate, abstracting from the source. I show that lower safe rates raise wealth inequality through two channels in the model: a heterogeneous impact on household wealth accumulation, and an incentive to take on riskier portfolio positions in search of returns.

Sections 3.3 and 3.4 together point to and quantify a distributional feedback loop in equilibrium, where exogenous increases to inequality (e.g. through technological change) lower safe interest rates, which then further increase wealth dispersion. As a result, the low interest rates and rising inequality seen today in advanced economies would be expected to persist, especially in light of the forecasts for further population aging and rising inequality across the globe.

3.2 why the lifecycle matters

3.2.1 *Introduction*

In this section, I analyze the component of household saving driven by lifecycle considerations. I do this by examining the effect of introducing lifecycle behavior to a canonical heterogeneous agent model.

Foundational papers in heterogeneous agent macroeconomics highlight the importance of the precautionary savings motive, showing how uninsured income risk at an individual level has aggregate effects, leading to an inefficiently high level of savings, implying over-accumulation of capital (Aiyagari 1994) and a depressed risk-free rate (Huggett 1993).

I find that adding lifecycle to those models exacerbates these effects. Lifecycle agents, facing a hump-shaped profile of expected income, will save when they work in order to finance consumption in retirement. How savings vary over the

lifecycle will also differ across income and wealth groups. It is therefore *a priori* ambiguous whether lifecycle motives increase or decrease aggregate household savings. I show that in a calibrated model, aggregate lifecycle savings are large and positive.

Lifecycle motives thus contribute to more over-saving relative to a complete market baseline. The need to smooth consumption against lifecycle patterns in income is another form of precautionary savings that exacerbates over-accumulation in these economies.

This section asks how large this effect is, quantitatively. Further, introducing a lifecycle will affect the savings behavior of rich and poor agents differently. How do savings change in the cross-section, and what are the implications for the distribution of wealth?

Intuition suggests that inequality between age groups should increase mechanically, a result of introducing more kinds of agents. However, the retirement savings motive turns out to decrease wealth inequality within age groups.

Below, I solve a version of the model of Section 1.2 disabling lifecycle effects, and compare it to the full specification, in order to tease out the *ceteris paribus* effect of lifecycle savings motives. I find that introducing the lifecycle increases aggregate savings and wealth in partial equilibrium by around 20%. The equilibrium interest rate correspondingly falls from 2.4% to 1%. This finding speaks to the risk-free rate puzzle (Weil 1989), which points out a theoretical difficulty in explaining the low risk-free rates prevailing in the data. Surprisingly, despite the large amount of extra heterogeneity in the form of age, overall wealth inequality is not much higher in the lifecycle model.

This section can be thought of as a complement to Chapter 4. There, I exam-

ined the effects of a small change in social insurance against lifecycle variation in income. This is equivalent to a marginal change in the household lifecycle saving motive. Here instead, I consider a larger discrete change in lifecycle behavior. As we shall see, many of the reactions are similar.

3.2.2 *Setup*

I now analyze the impact of introducing a lifecycle to a stationary model. I do this by solving a version of the model from Section 1.2 in which I kill off all lifecycle effects, and comparing it to the full model. This means that instead of an age profile y_t of labor income, agents receive a constant wage y multiplying their productivity. Lifespans remain uncertain: I switch to a perpetual youth setup where the flow mortality probability is the same at every age, $\lambda_t \equiv \lambda$. I set y to the average of the profile y_t calibrated above, weighted by the age structure, and the mortality rate λ is chosen to match life expectancy at age 25 in the calibrated model,

$$\frac{1}{\lambda} = \int_0^\infty e^{-\int_0^t \lambda_s \, \mathrm{d}s} \, \mathrm{d}t \approx 55. \tag{3.1}$$

I use perpetual youth rather than infinitely lived agents in this experiment for two reasons. Firstly, perpetual youth models do not feature lifecycle in any real sense: households of every age have the same policy functions. Age only matters insofar as it gives older agents more time potentially to accumulate wealth. There is no notion of households smoothing consumption against lifecycle variation in income.

Secondly, perpetual youth allows us to retain comparability to the lifecycle model. There, mortality risk is a factor in households' consumption-savings decision that would be absent with infinitely lived agents.

I make a few other changes in this section relative to the calibration of Section 1.3. First, I turn off the risky asset k in both the lifecycle and perpetual youth economies. This means that initial conditions and income shocks are the only source of inequality within age groups. The equilibrium rate r^b adjusts to keep aggregate household wealth equal to government borrowing. By shutting off the exogenous returns coming from k, we can read more meaning into r^b as the interest rate that balances all household borrowing and lending, providing an additional tool for comparing the lifecycle and perpetual youth economies.

Second, I impose a fixed exogenous level of government debt in both economies, set to U.S. public debt per capita in 2016 ($60,000). Otherwise, aggregate household wealth in the model would have to be zero in equilibrium, which can imply a negative equilibrium interest rate.[1]

Finally, fiscal policy in the full model redistributes wealth across income groups and thus implicitly across ages. This can mask the full impact of lifecycle variation in income. I minimize this effect by disabling fiscal policy as much as possible in both settings.[2]

3.2.3 *Introducing lifecycle*

Table 3.1 lists important aggregate moments in the lifecycle and perpetual youth economies. Saving for retirement leads aggregate wealth and savings to be significantly higher in the lifecycle model. Fixing the interest rate (column 2), we

[1] Negative interest rates cause technical difficulties, especially with boundary conditions.

[2] Seniors in the lifecycle model still need to receive a positive income, in order to be able to consume *something* at every wealth level. As a result, I maintain small social security transfers paid for by a linear income tax in the lifecycle model. This fiscal policy slightly smooths out lifecycle variation in income, so that the results of this section may understate the true effects of lifecycle motives.

Table 3.1: Effect of lifecycle: model aggregates and inequality. Aggregates in thousands of dollars.

	Lifecycle	Non-lifecycle	
		Partial eq.	General eq.
Interest rate	1.0%	1.0%	2.4%
Agg. household wealth	60	49	60
Agg. consumption	56.0	56.1	56.5
Agg. savings	1.6	1.4	2.0
Income Gini	0.90	0.85	0.85
Within-age wealth ineq.*	2.26	2.31	2.20
Between-age wealth ineq.*	0.16	0.03	0.05

* Measured using Theil coefficients

see that aggregate savings and wealth are both 15-20% higher in the lifecycle model. As a result, the market clearing interest rate falls from 2.4% to 1% in the lifecycle model. The over-saving effects documented by Aiyagari (1994) and Huggett (1993) are exacerbated.

Income inequality is higher in the lifecycle model due to age variation in the wage. Wealth inequality overall is also higher in the lifecycle model. The table reports wealth inequality using the Theil index of inequality, a measure that lets

us decompose wealth inequality into inequality between and within age groups.[3]

There is very little between-cohort wealth inequality in the perpetual youth model. The only source of variation in wealth by age in that case comes from deterministic growth in wealth, which occurs in when $r^b - \rho + \lambda - \pi \neq 0$. This depends on r^b, but is quite small compared to wealth fluctuations that stem from lifecycle savings motives.

By contrast – and rather surprisingly – disparities within age groups are higher in the perpetual youth model. We can understand this by looking at Figure 3.1, which plots the Theil *within-* coefficient of inequality for each age group in both economies, together with average savings. The two plots mirror each other: wealth inequality tends increase in periods of the lifecycle where agents want to save.

This can be explained by the presence of constrained agents, who are consuming their income hand-to-mouth and saving almost zero. There are hand-to-mouth agents of every age. In periods of the lifecycle when agents tend to want to save (during working life), those who have enough for a baseline level of con-

[3]For a probability measure $\mu(\mathrm{d}x)$ with positive support, the Theil index is

$$\int_0^\infty \frac{x}{\mathbb{E}_\mu(X)} \log\left(\frac{x}{\mathbb{E}_\mu(X)}\right) \mu(\mathrm{d}x). \tag{3.2}$$

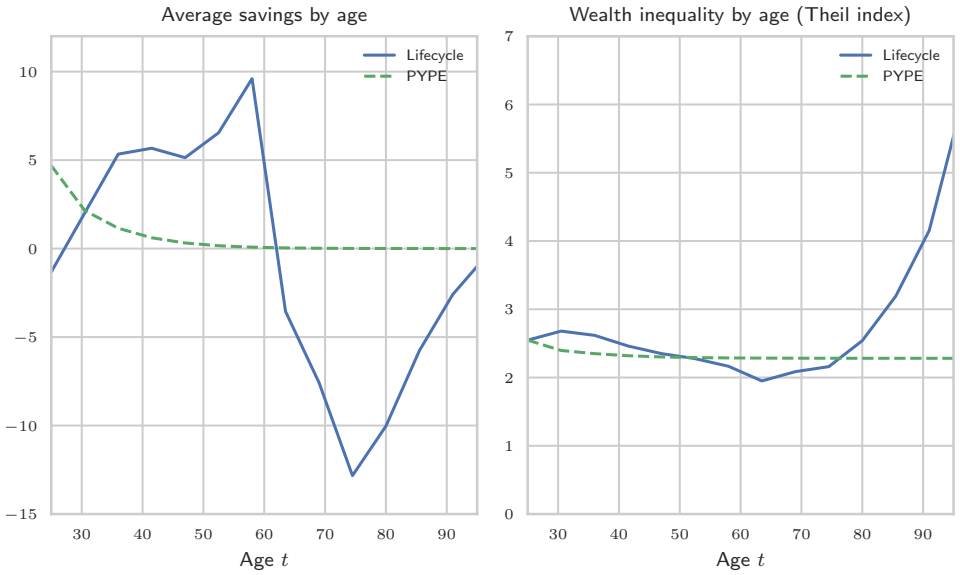
It is decomposable as follows. Denote by $\mathcal{A}$ the set of agents. For any $A \subseteq \mathcal{A}$ denote by $T(A) = \int \frac{w}{\mathbb{E}(W|A)} \log\left(\frac{w}{\mathbb{E}(W|A)}\right) m(\mathrm{d}w|A)$ the Theil index for wealth of agents in A, where $m(\mathrm{d}w|A)$ is the marginal probability of w conditional on A. Further, denote by $s(A) = P(A)\frac{\mathbb{E}(W|A)}{\mathbb{E}(W)}$ the share of total wealth held by agents in A. If Π is a partition of $\mathcal{A}$ into disjoint subsets, then

$$T(\mathcal{A}) = \underbrace{\sum_{A\in\Pi} s(A)T(A)}_{\text{within}} + \underbrace{\sum_{A\in\Pi} s(A)\log\left(\frac{E(W|A)}{E(W)}\right)}_{\text{between}}. \tag{3.3}$$

The first term is an average of inequality inside each subgroup, weighted by their share of wealth, measuring within-subgroup inequality. The second term measures inequality in the subgroup means, capturing inequality between subgroups.

Figure 3.1: Savings and wealth inequality in lifecycle and perpetual youth models

Note: Average savings and wealth inequality plotted by age group in the lifecycle and perpetual youth economies (fixing r^b – partial equilibrium). Savings in thousands of dollars per year.



sumption today will increase savings and accumulate more wealth. Constrained agents, however, cannot do so adequately and have high marginal propensities to consume. As a result, better off agents accumulate more wealth while the poorest cannot, leading to an increase in wealth inequality. Conversely, in periods of dissaving (retirement), most agents decumulate wealth but the poorest are hand-to-mouth and cannot, tightening the wealth distribution.

The overall impact of lifecycle patterns on aggregate wealth inequality thus depends on the relative masses of savers, dissavers, and hand-to-mouth agents over the lifecycle. Quantitatively, this sums up to a *negative*: lifecycle motives lead to an increase in within-cohort inequality.

However, overall wealth inequality increases. The decrease in lower within-

cohort inequality is dominated by the higher inequality between cohorts, stemming from lifecycle fluctuations in income.

This change in the makeup of wealth inequality is important. For example, through self-insurance, the more income varies over the lifecycle, the lower is wealth inequality among seniors.

As for the general equilibrium response, the risk-free rate r^b falls by 1.4 percentage points when lifecycle is introduced. As a result, we see an increase in aggregate household savings and consumption, together with a rise in wealth inequality. These mechanisms will be discussed in detail in Section 3.4, where I analyze the *ceteris paribus* effects of a change in the risk-free rate.

3.2.4 *Conclusion*

This section strives to quantify the effect of lifecycle motives on savings behavior. I have shown, using an overlapping generations heterogeneous agent model, that lifecycle behavior has a large effect on aggregate variables, even once the precautionary saving motive has been taken into account.

This experiment generalizes Aiyagari (1994) and Huggett (1993), who stressed the importance of labor income risk to aggregates. These findings help explain the increasingly low interest rates prevailing in advanced economies. Population aging widens the variation in lifecycle income, and can thus lead to more aggregate saving.

This section can be thought of as a discrete analog to the marginal change in lifecycle behavior studied in Chapter 4. There, we shall find similarly that a small increase in social insurance against lifecycle variation (which is equivalent to a small decrease in the lifecycle savings motive) associated with higher inequality

and lower aggregate savings, in line with the results of this section.

Extensions should verify some of the elasticities derived in this section against international evidence, looking at the relationships between lifecycle variation in income, aggregate savings, and wealth inequality.

3.3 population aging and rising inequality

Population aging and rising inequality are two global trends that share several characteristics, as argued in the thesis introduction. This section seeks to quantify their macroeconomic and distributional effects and analyze the relevant channels, by harnessing the framework of Chapter 1 to treat these trends in a symmetric way. Section 3.3.1 analyzes the impact of population aging, looking at an increase in life expectancy as observed in Japan in the past decades, and evaluates fiscal policy reactions. Section 3.3.2 looks at the impact of rising income inequality due to exogenous factors. In both cases, household savings increase, leading to a decline in the equilibrium interest rate. The distributional effects are mixed: I find that population aging reduces wealth inequality somewhat, while rising income inequality increases that measure. Regardless, the decline in safe interest rates itself tends to disperse the distribution of wealth, an idea I explore in Section 3.4. The question remains also of the appropriate policy response to these distributional changes, a topic I pick up in Chapter 4.

3.3.1 *Demographic change*

The solid lines in Figure 1.2 plot the evolution of the age structure in Japan, strikingly illustrating the aging of the Japanese population, which results from

lengthening life expectancy and falling birth rates. Japan is not alone in this, United Nations (2015) forecasts that the world dependency ratio (ratio of seniors to people of working age) will increase threefold between 2015 and 2100. In this section, I focus on the example of Japan – with its most aged population, it has been argued to be a harbinger of things to come for all advanced economies.

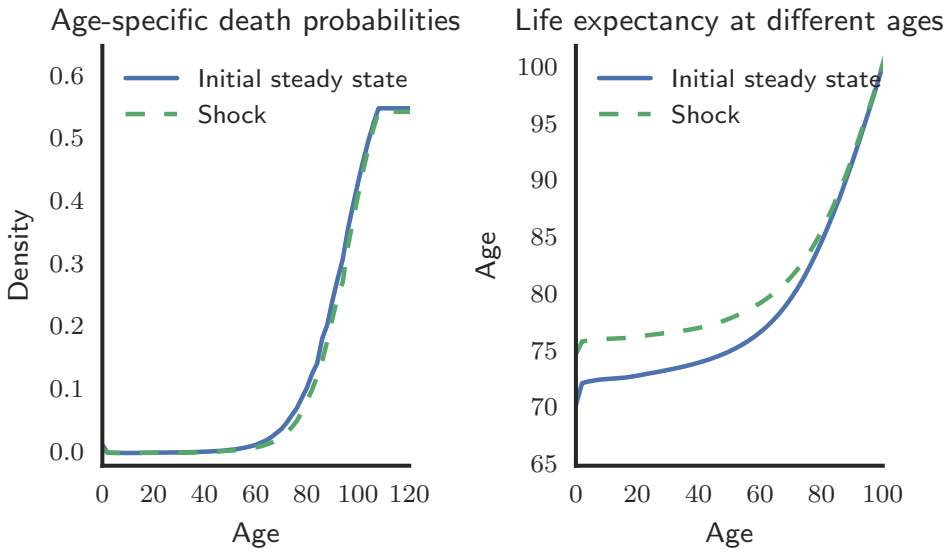
Population aging affects the macroeconomy and asset prices through a number of channels: the lifecycle savings motive is amplified, the fraction of retirees increases, capital per worker increases, labor supply per capita falls, and so on. Carvalho et al. (2016) argue that of these, those with the largest effect on the natural rate of interest are the altering of the household's consumption–savings decision, and the compositional channel. I focus on these two in this section.

Further, there are a number of important fiscal implications, most notably on the social security system. For example, International Monetary Fund (2017) predicts a 35% increase in real pension spending and a 118% increase in real health care costs in the United States (which is not so advanced in the demographic transition) between 2010 and 2050. Countries face an array of policy choices for dealing with a social security shortfall. For example, they can cut costs by decreasing benefits or extending the retirement age. They can raise taxes in a number of ways, or defer taxation by issuing debt in the short term.

The model is well placed to analyze these different policy responses to population aging. This section presents a test bed in which to look at these policies, focusing on the transition dynamics resulting from lengthening lifespans in the framework of Chapter 1.

Figure 3.2: Population aging: the demographic shock

Note: The shock decreases age-specific mortality probabilities from 1970 to 1980 levels. Life expectancy increases at each age correspondingly.



Population aging

In the transition equilibrium studied in this section, there is only one shock: age-specific mortality probabilities fall instantly from levels observed in Japan in 1970 to those of 1980, plotted in Figure 3.2. (Birth rates are kept fixed.) This equilibrium will serve as a baseline for policy experiments.

The baseline model follows as far as possible the calibration set out in detail in Chapter 1. The demographic parameters are calibrated on Japan, using the All Japan Life Tables computed by the Japanese Mortality Database (Japanese National Institute of Population and Social Security Research 2017).The economy begins with the age-specific mortality probabilities prevalent in 1970. The birth rate is set so that total population does not change in the initial steady state.

In order to start from a demographic steady state, I start the economy with a stationary demographic distribution. This is the age structure derived from freezing the mortality probabilities in time, assuming perfect replacement, and letting time run to infinity. It is plotted as the dashed line in Figure 1.2, and is close to being an average of the empirical age structures of the past decades.

At time $\tau = 0$ in the transition, all the age-specific mortality probabilities jump from 1970 to 1980 levels, and the demographic distribution evolves naturally. Birth rates are fixed at their initial level – the aim of this section is to analyze the impact of lengthening lifespans alone.

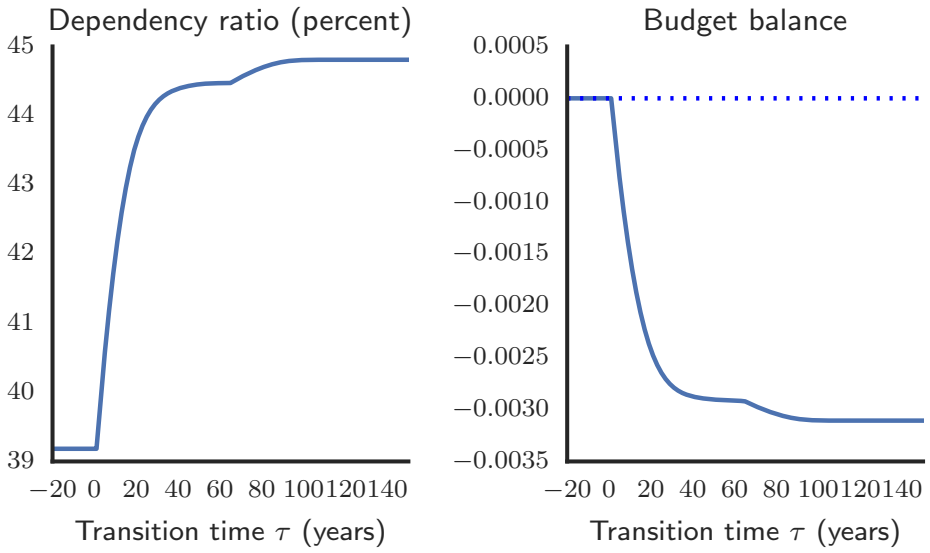
The government provides a flat social security payout to agents aged 65 or above. This is financed by taxing the working age population. The social security payout is fixed at \$10,000/year, and the taxes are calculated so that the government is in budget balance before the shock hits. Government consumption and debt are zero, and no other taxes are in place at this point. The extra social security liabilities resulting from population aging will be financed through an increase in the linear income tax.

As the shock hits and agents begin to live longer, the dependency ratio eventually increases from 39% to 45% (pictured in Figure 3.3). Empirically, the Japanese dependency ratio is projected to continue to increase far past this point (Economic Co-Operation and Development 2015) – recall that we are only considering a one-time increase in life expectancy in this experiment. If the government maintains its original fiscal stance, this leads to a budget deficit in each period, pictured in the right panel of the figure.

How should the government finance this shortfall? Before addressing this question, I examine the results of the model in the case where the shortfall is not

Figure 3.3: Population aging: Dependency ratio and budget pressure from population aging in transition

Note: The dependency ratio (fraction of elderly relative to workers) increases (left), leading to a budget shortfall (right) in the absence of further policy intervention. RHS units: \$100,000/capita

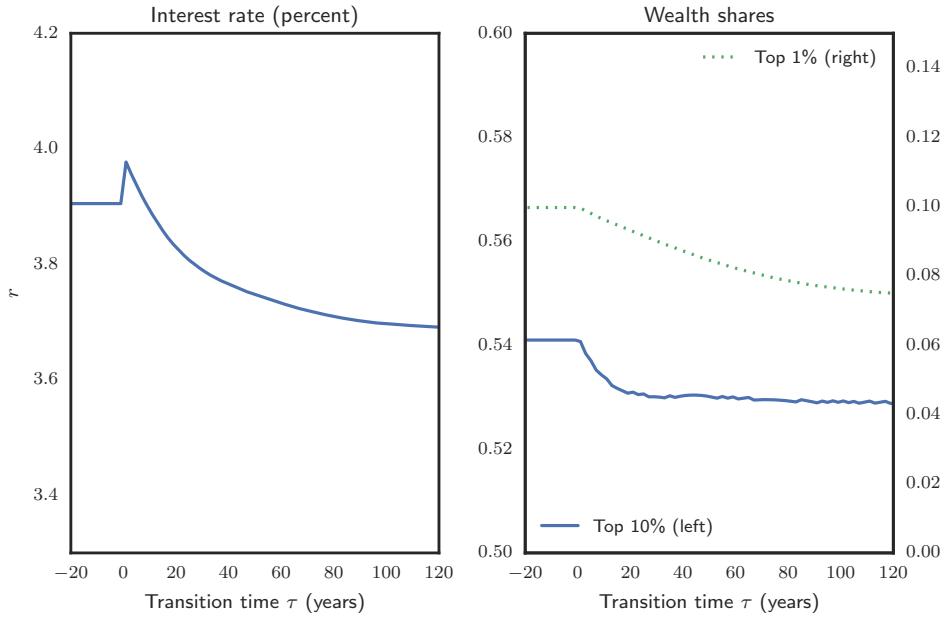


financed, as it serves as a useful baseline against which to compare the policy responses.[4]

Figure 3.4 plots the reaction of the interest rate over time to the mortality shock, as well as a measure of inequality. The interest rate moves by around 30 basis points. It initially jumps upwards, a consequence of the surprise increase in permanent income due a longer anticipated retirement period receiving benefits. It then gradually declines to its lower steady state level, as lifecycle savings increase.

[4]Note that strictly speaking this baseline case is a partial equilibrium: for the purposes of this experiment I assume the government finances this shortfall out of pocket without participating in the bond market.

FIGURE 3.4: POPULATION AGING: EQUILIBRIUM INTEREST RATE AND WEALTH INEQUALITY IN TRANSITION



Although this is not a large movement in interest rates relative to that seen in the data, it is important to remember that this is a plot of one decade's worth of improvement to life expectancy. In the data, there is a cumulation of many persistent improvements of this nature, as well as non-demographic factors leading to the significant decline in interest rates observed in Japan.

From the right-hand panel of Figure 3.4, we see that inequality is dropping in response to the demographic shock. The top 10% wealth share drops relatively fast by around 1 percentage point. Interestingly, the top 1% share drops by around 2 percentage points, a much larger adjustment proportionally, and does so more gradually. With a longer anticipated retirement, households face strong pressure to increase their savings. Poorer households sacrifice working-age consumption to build up their wealth as a result. By contrast, those households

with ample wealth and sufficient retirement savings need only to adjust somewhat their wealth accumulation. Thus, the distribution of wealth compresses somewhat.

Define welfare in the utilitarian sense as the population-weighted average value function of households,

$$W_\tau = \int V(\tau, t, w, z)\, \mathrm{d}\mu.$$

Welfare naturally increases as agents live longer, consuming for longer and receiving the social security benefits of a longer retirement. As the interest rate declines, this has a negative wealth effect on agents who are lenders, but a positive effect on the new young. These effects balance out to a net increase in welfare over time.

Policy experiments

This section evaluates different policy options for financing the social security shortfall and compares them to the baseline model from the last section. These policy options fit into two categories: tax increases and cuts to benefits.

I begin by discussing defined benefits schemes, in which the structure of benefits to the elderly is fixed and the government instead increase taxes. I assume the government keeps budget balance at every period, and consider a flat lump-sum tax levied on all agents of working age, and a linear income tax increase.

I also evaluate one defined-contribution scheme, in which benefits are cut in the form of an increase in the retirement age. The government endogenously chooses the retirement age in order to avoid a budget deficit in any period.[5]

[5]Note that the retirement age is assumed to be an integer, and the government makes small surpluses from the rounding difference. I assume this is consumed by the government.

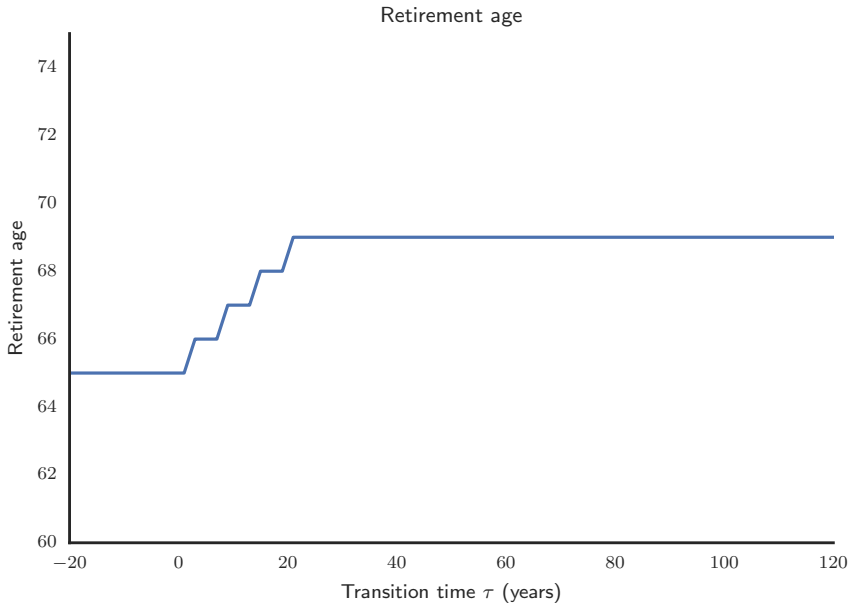
In response to this improvement in life expectancy equivalent to ten years of medical progress, the path for the retirement age that maintains budget balance, pictured in Figure 3.5, gradually takes it from 65 to 70 over the course of twenty years.

The results of these policy experiments are pictured in Figures 3.6 to 3.8. I find that increasing the retirement age leads to the best welfare outcome in the long run, although for the first 20 years after the shock this is dominated by an increase in the income tax. Myopic policymakers with short mandates might then choose to continue to postpone structural changes such as increasing the retirement age and raise taxes instead.

Raising the retirement age also has a marked effect in decreasing wealth inequality. This can be understood because the rich in the model are generally old. A cut to retirement benefits can be thought of as a tax on the elderly, so that raising the retirement age acts like a tax with relatively high incidence on the wealthy. Further, it raises the share of labor in total household income, which is generally less dispersed than capital income. These channels lead to a decrease in wealth inequality.

Another surprising result is that the lump-sum tax leads to less inequality than the income tax. This can be explained by the lifecycle patterns in wealth and savings: it is precisely when they are earning the most that agents save the most, and so taxes at that time are especially painful in terms of lifetime income. Because households tend to be borrowing constrained in the beginning of their working life, they are not able to smooth the tax incidence adequately. The lump-sum tax, by spreading out the incidence over their working life, provides a tax smoothing vehicle for them which leads to better wealth outcomes. Indeed,

Figure 3.5: Population aging: The endogenous increase in retirement age required to maintain budget surplus



as shown in Figure 3.9, total wealth is higher under the lump-sum tax scheme.

Note however that aggregate welfare is better under the income tax than the lump-sum tax, since the lump-sum tax forces constrained agents into lower consumption.

To sum up, we see that population aging depresses the natural rate by raising aggregate savings through two channels: households save more for retirement over the lifecycle, and the population share of seniors – who tend to have relatively high wealth – increases. Wealth inequality falls somewhat due to stronger substitution among poorer households. Of the policy options for financing a social security shortfall in the model, raising the retirement age has the best long-run welfare outcomes, but a short-sighted policymaker would prefer to make a nonstructural adjustment such as raising taxes.

Figure 3.6: Population aging: The equilibrium path of the interest rate under different policies

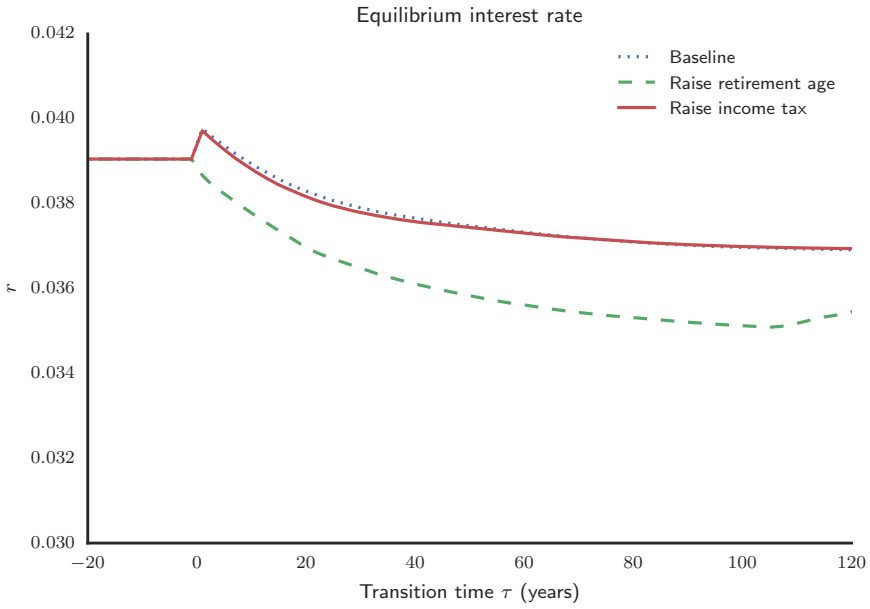


Figure 3.7: Population aging: The path of the shares of wealth held by the top 10% and 1% for different policies

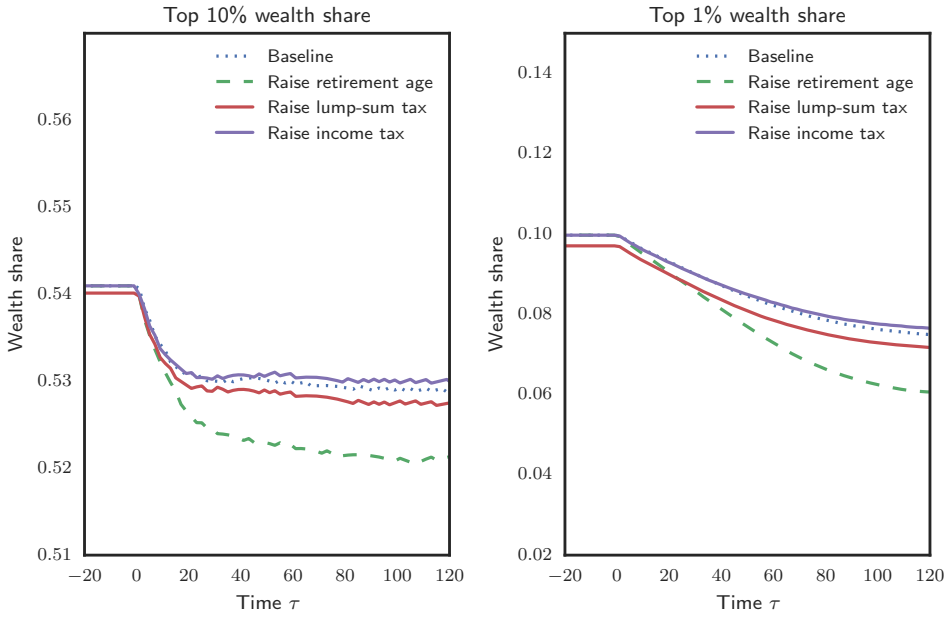


Figure 3.8: Population aging: Aggregate welfare over time for different policy responses

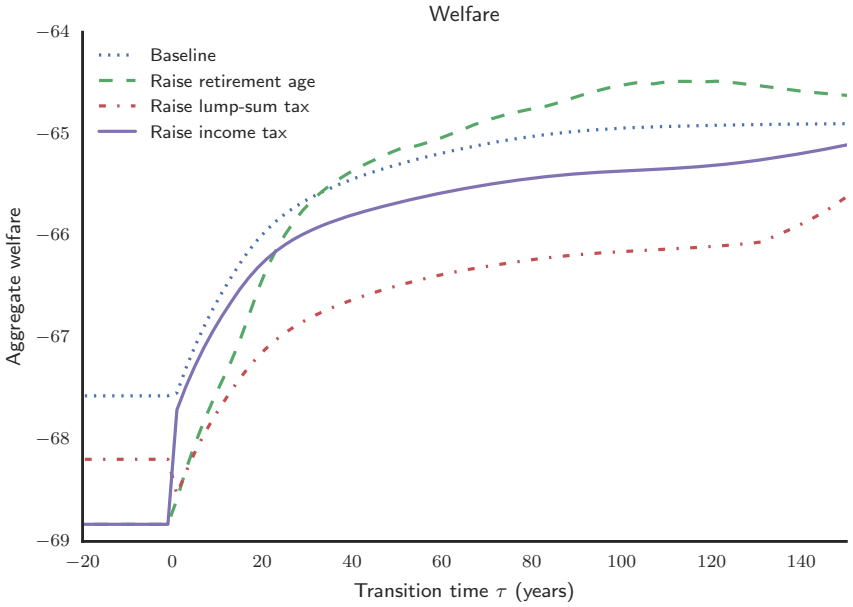
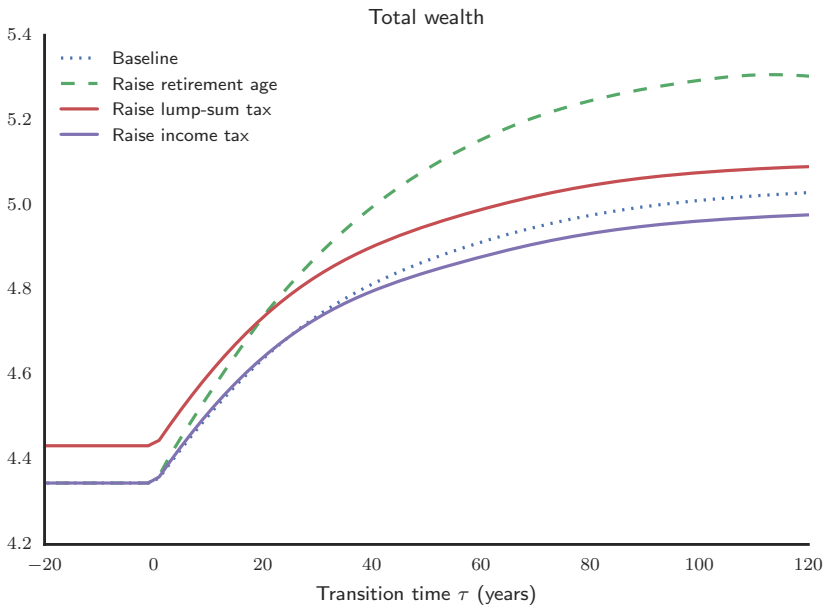


Figure 3.9: Population aging: Aggregate household wealth under different policy responses



3.3.2 *Rising inequality*

Labor income inequality has been increasing in the United States over the past decades. One explanation for this phenomenon is skills-biased technological change, such as the rise of artificial intelligence, which increases the productivity of high-skilled labor. The persistence of this trend has placed renewed emphasis on understanding the implications of changes to the distribution of household income. Keynes (1936), Dynan et al. (2004), and Summers (2015) argue that rising income inequality is associated with higher savings rates. In this section, I examine the implications of rising income inequality in the model, and find support for this hypothesis. Nevertheless, the empirical evidence on this relationship is somewhat mixed (Bertrand and Morse 2016; Bofinger and Scheuermeyer 2016). Heathcote et al. (2010) examine the welfare implications of this increase, using a model more sophisticated in some axes (education choice, labor choice, intra-household decisions) and more stylized in others (perpetual youth).

The broader macroeconomic and distributional impact of a change in the distribution of household income can be additively decomposed into two parts. The first is a shift in the mean, i.e. of macroeconomic labor productivity. The macroeconomic effects of changes in aggregate labor productivity have been well studied. When aggregate productivity grows, this is associated with higher household income, which should lead to more savings and lower equilibrium safe rates.

The second component, which I focus on in this section, is a change in the distribution of income that leaves the mean unchanged. First order is the *compositional channel*: richer households tend to have higher savings rates, so that a mean-preserving increase in income inequality should raise aggregate wealth and

depress safe interest rates. The aim of this section is to quantify these effects using the framework of Chapter 1.

A mean-preserving increase in income inequality can be modelled in a number of ways in this setting, which could capture different aspects of the changes observed in the empirical income distribution. One would be an increase in the size or frequency of the idiosyncratic shocks to labor income. Beyond the compositional channel, more uncertainty would lead households to hold larger precautionary savings, further depressing the equilibrium rate.

Instead, I focus on the productivity of newly formed households, increasing the inequality of the distribution from which it is drawn. This could result from increasing inequality in educational attainment or human capital. Indeed Huggett et al. (2011) and Storesletten et al. (2004) argue that initial conditions are empirically even more important than shocks over working life for explaining the variation in realized earnings and wealth. This modelling approach has the advantage of not directly altering policy functions, thus isolating the compositional channel.

Another alternative would be to widen the distribution of household capital income, by increasing the volatility of the idiosyncratically risky asset which represents returns to private housing and private business investment. Indeed, the labor share of household income has been declining in many countries (Piketty et al. 2017). This would yield a parallel channel for the increase in household income inequality alongside the diffusion of the distribution of labor income. In the model, this would affect the household portfolio decision, somewhat muddying the effect of the compositional channel.

The calibrated distribution of initial productivity is widened with a mean-

preserving spread, so that the Gini coefficient this distribution increases from 0.47 in the baseline calibration to 0.56. In response, the long-run equilibrium interest rate falls from 2% to 1.9%. This corresponds to a -26% elasticity of the neutral rate to the Gini coefficient for initial productivity.

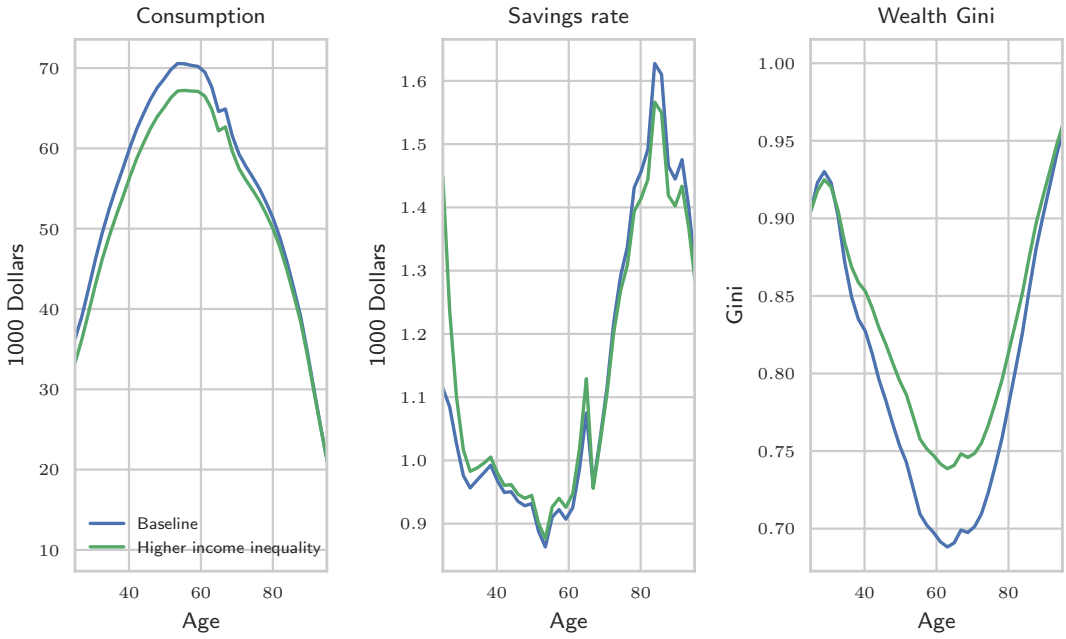
Figure 3.10 plots average consumption and savings rates over the lifecycle, alongside age-specific wealth inequality, for both the baseline case and that with higher income inequality. We see that consumption falls and savings rise, despite the fall in the interest rate. This is due to the shift in the distribution of wealth towards wealthier households, which have higher savings rates as discussed in Chapters 1 and 4.

The diffusion of the income distribution leads to more wealth inequality over the lifecycle. This is most pronounced for households around retirement age, at the apex of lifecycle wealth accumulation. The Gini coefficient for the overall wealth distribution increases by 3.5 percentage points in general equilibrium. Note that a fall in the risk-free rate will itself lead to higher wealth inequality by incentivizing riskier investment, an idea I explore in more detail in the next section. Nevertheless, even in a partial equilibrium fixing the risk-free rate, the wealth Gini coefficient increases by 3 percentage points, so that the bulk of the adjustment can be ascribed to the compositional channel, in which higher income inequality places more income into the hands of households with higher savings rate by virtue of their large wealth.

We thus find a relatively large negative elasticity of the risk-free rate to the structural income inequality. Note that in this section I focused only on the compositional channel. In practice, different sources for widening inequality could further increase household savings and amplify the reaction of the risk-free rate.

Figure 3.10: Effect of higher income inequality on lifecycle consumption–savings and inequality

Note: Figure plots long-term profiles of mean consumption and savings rate by age, as well as the Gini coefficient of the age-specific wealth distribution. These are plotted both in the baseline specification, and in the case with more diffuse distribution of household labor productivity.



3.3.3 *Conclusion*

To sum up, we see that population aging and rising income inequality both increase aggregate savings through a number of channels, thus depressed safe interest rates. Especially important is the *compositional channel,* where more wealth is held in the hands of agents with higher saving rates. The distributional impact of these trends differs, with population aging leading to a slightly more concentrated wealth distribution in the model, while rising income inequality tends to disperse wealth. The fall in the equilibrium interest rate itself has a number of

distributional effects, which we turn to in the next section.

This section focused on quantifying the compositional channel, highlighting the versatility of the framework of Chapter 1. Extensions of this work to study secular stagnation more broadly could consider additional mechanisms, such as labor choice, supply-side effects, and varying risk premia.

3.4 DISTRIBUTIONAL EFFECT OF A PERMANENT INTEREST RATE CHANGE

3.4.1 *Introduction*

In this section, I examine the distributional impact of a permanent exogenous change in the equilibrium risk-free rate r^b. This experiment elucidates the distributional mechanisms of the model in Chapter 1. Indeed, these analyses are relevant to all chapters in this thesis, since each one entails a general equilibrium adjustment of r^b.

In the last section, we discussed the idea that higher inequality depresses the natural rate of interest by raising total household saving, since well-off households tend to have higher savings rates. In this section, I examine the other causal direction: asking what are the distributional effects of a change in the natural rate, regardless of its origin. I find that in the model, a fall in the interest rate leads to higher inequality through a number of channels. If the interest rate fell due to an exogenous increase in inequality (such as structural changes in the labor market), these mechanisms point to an amplification of the distributional shock in general equilibrium. If the fall is due to other factors (such as population aging, as analyzed in Section 3.3), this analysis helps predict the resulting changes in the distributions of household wealth and consumption.

In response to an decline in the risk-free rate, the model predicts a long-run increase in consumption, savings, and household wealth. The richer and older the household, the stronger the effect. As a result, wealth inequality within age groups rises, especially so among seniors.

This reaction can be explained by patterns in household bond portfolios. An overwhelming majority of households are indebted, using leverage to maximize their investment in the high-yielding risky asset. The lower cost of borrowing raises their total return on investment, and thus their disposable income.

Typically, after a fall in the interest rate, borrowers gain and lenders lose out. However, lenders in the model tend to be more aged and more wealthy. They were also usually net borrowers of bonds earlier in their life. As a result, their loss is outweighed by the gains from paying lower borrowing costs in their youth.

When safe interest rates fall, households buy more of the risky asset. Since this asset carries idiosyncratic risk, this leads to an increase in wealth inequality. This is the principal mechanism linking low safe interest rates to higher inequality discussed in the literature.[6]

This section relates to the growing literature studying the interactions between demographics, distributions, and monetary policy, including Auerbach and Kotlikoff (1985), Wong (2015), Auclert (2017), and Sterk and Tenreyro (2018).

This section focuses on the direct effects of the change in the interest rate on households. Indirect effects deriving from general equilibrium reactions, such as changes to wages and transfers, should be analyzed in an extension to this work. Evidence on the distributional effects of monetary policy Kaplan et al. (2018)

[6]See, for example, Piketty and Saez (2003) and Saez and Zucman (2016). Kuhn et al. (2017) has a great analysis of the impact of asset price changes on the wealth distribution in the U.S. since 1949.

and Ampudia et al. (2018) suggests that – at least in the short run – indirect effects account for a large portion of monetary transmission.

3.4.2 *Experiment*

I take the economy calibrated as in Section 1.3, and permanently increase government consumption from 0 to \$30,000 per capita, financed by extra government debt. As a result, the market-clearing rate for bonds increases from 2% to 3%.This is a simple way of implementing a ceteris paribus change in the interest rate. In practice, these could result from changes in the age structure or labor market (as in the last section), or to fiscal, monetary, or financial factors.

I focus below on an increase in the interest rate, but the results below can be carried over to a fall in r^b by reversing the sign (numerical experiments suggest that reactions are relatively linear in size and sign of the interest change, away from the zero lower bound).

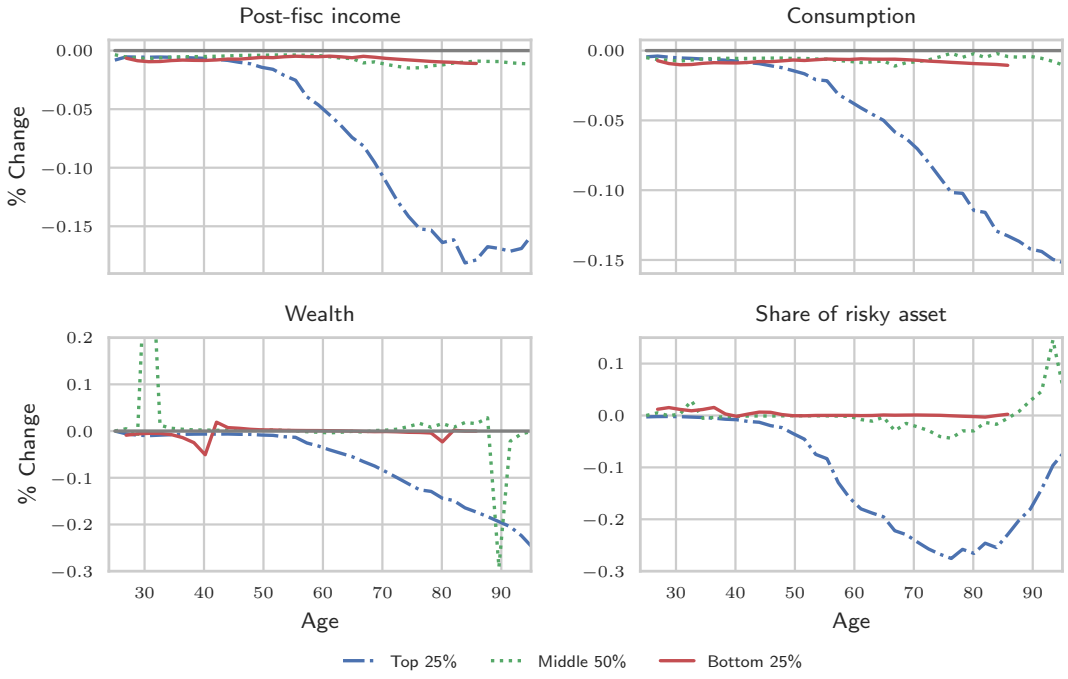
Figure 3.11 plots comparative statics for three different wealth groups, the difference between lifecycle profiles in the low-r^b and high-r^b steady states. Disposable income, consumption, savings, and household wealth drop for all households as r^b increases. The richer and older the household, the stronger the effect.

The key to understanding this reaction is the portfolio decision of households. The two left-hand panels of Figure 3.12 illustrate that lenders in the economy are overwhelmingly older and towards the top of the wealth distribution. The most households hold net short positions in bonds, leveraging into the high-yielding risky asset, despite positive net worth. The income effect due to higher repayment costs reduces consumption and savings across the board.

This is felt especially strongly by wealthier households, with a high capital

Figure 3.11: Effect of interest rate hike on lifecycle profiles

Note: Figure plots change in long-term lifecycle profiles in response to 1 percentage point increase in interest rate r^b, by wealth group.



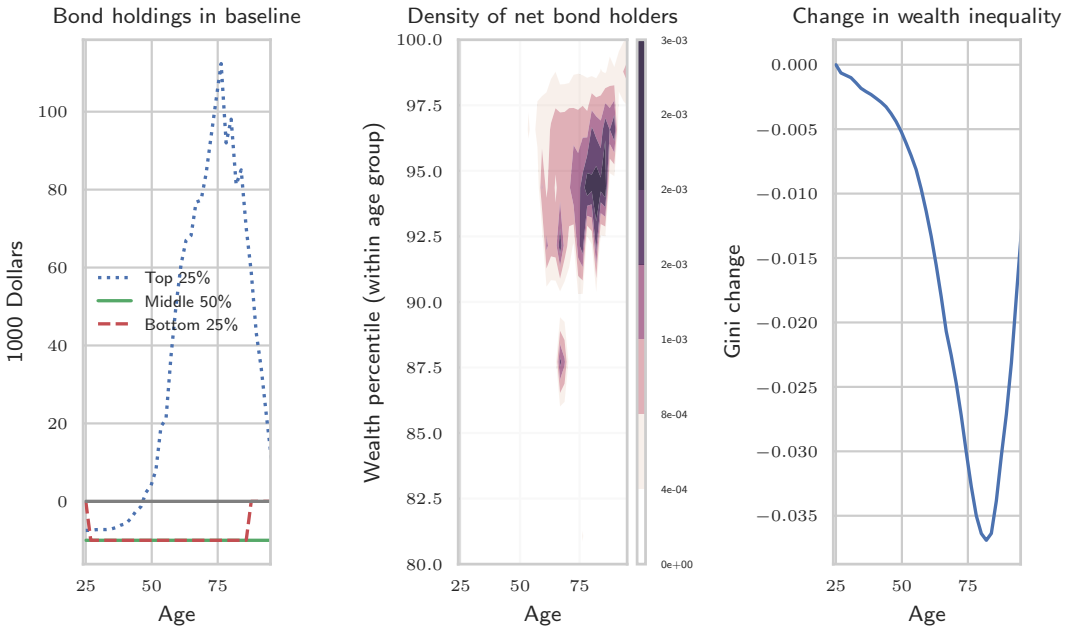
share of income. Since wealth – and thus capital income – is determined by cumulated savings, the reaction strengthens through the lifecycle of these households, as visible in Figure 3.11.

Wealth inequality within age groups declines as a result, as evident in the right-hand panel of Figure 3.12. The compounded effect of lower returns reduces the wealth of top asset holders more than lower wealth groups, compressing inequality.

There are two other important ways in which the increase in the risk-free rate affects the distribution of wealth in the model. These have to do with households

FIGURE 3.12: BOND HOLDINGS AND WEALTH INEQUALITY AFTER AN INTEREST RATE HIKE

Note: Figure plots lifecycle bond holdings in the baseline (left), the density of those with net positive bond holdings (middle), and change in wealth inequality within age groups (right) in response to 1 percentage point increase in interest rate r^b.



adjusting their portfolios.

The bottom right panel of Figure 3.11 shows that when r^b increases, households tend to find bonds more attractive and rebalance in their favor, away from risky capital. By simple virtue that they are holding fewer assets with idiosyncratic risk, wealth inequality decreases. This is the standard *flight-to-yield* argument linking low safe rates with higher wealth inequality.

Indeed, we can see this in the model by rewriting the first-order condition for capital (Proposition 1.1) as

$$\frac{k}{w} = \frac{1}{\epsilon}\frac{1}{\gamma}\frac{r^k - r^b}{\sigma^2}, \tag{3.4}$$

where

$$\epsilon = \frac{\partial_w c}{c/w}$$

is the elasticity of consumption to financial wealth.[7] The optimal portfolio share of capital in financial wealth is given by the risk premium discounted twice: by the coefficient of relative risk aversion and by this elasticity. The more fluctuations in wealth (and thus in k) pass through to consumption, the more agents dislike risk and shy away from k.

This flight-to-yield mechanism simply comes from a decrease in the numerator of the right-hand side of equation (3.4). However, another channel that emerges from the model is a change in effective risk aversion, through ϵ. Higher borrowing costs reduce the return on wealth and lead households to accumulate less wealth. Households in the model are more risk-loving at lower levels of financial wealth; especially so for retirees.

This may sound paradoxical but is in fact optimal. Consumption is determined by total wealth, financial and human. For richer retirees who have little remaining permanent income, financial wealth w is the most important factor. The elasticity ϵ for these agents is close to one. By contrast, households that are poorer and younger face borrowing constraints and anticipate future income. Their elasticity of consumption to wealth is below one. In other words, safe in the knowledge that they have other sources of income, poorer and younger households are relatively more willing to accept risk.

As a result, households have an incentive to invest further into risky capital through this channel. However, numerically this channel is dominated by the

[7] Substitute the first order condition for consumption (and its derivative) into the condition for capital.

increase in the numerator, and overall households choose to deleverage.

Estimates of these reactions in the data vary. Ameriks and Zeldes (2004) do not find evidence of age variation in risky portfolio shares, despite the folk wisdom on personal investing over the lifecycle which advises gradual de-risking of the portfolio. Blundell et al. (1994) estimate that the intertemporal elasticity of substitution increases with age and income, in contrast with the predictions of my model. This can be ascribed to the simple household preference structure; extending the model to include non-homotheticities such as age-dependent elasticity of substitution would help match this fact (see Straub (2017) for a discussion).

Finally, it is important to discuss the lenders. As discussed, and charted in the middle panel of Figure 3.12, lenders in the model tend to be seniors, and towards the top of the wealth distribution. Lenders typically stand to gain from an interest rate rise. Here, however, even they are worse off in the long run. These wealthy seniors have had to go through a lifecycle in which they will have spent much of their lives as borrowers. As a result, they will have paid the higher borrowing rates and suffered the costs of the rate rise. This negative scarring effect dominates the larger returns that they now earn. Thus their wealth, consumption, and savings are below the baseline levels.

This thinking illustrates the importance of considering the lifecycle alongside heterogeneity in income and wealth. An analysis of an interest rate change, which looks at the impact on borrowers and lenders without acknowledging that most savers were once borrowers, would miss this channel.

3.4.3 *Conclusion*

This section examines the long-term distributional impact of a permanent change to safe interest rates. The classical decomposition of the winners and losers from an interest increase into lenders and borrowers is blurred in a model with overlapping generations and household heterogeneity in income and wealth. Here, as in the data, net lenders tend overwhelmingly to be richer and older. However, even they were often borrowers in their youth. Lifecycle and cross-sectional patterns in bond holdings imply that lenders are typically older and richer, and generally many of them were borrowers in their youth. As a result, even those who end up as lenders will have paid higher borrowing cost, and end up on average worse off than in the baseline experiment.

This experiment highlights that models with heterogeneity only in income and wealth are not enough for certain distributional questions in macroeconomics. Omission of the lifecycle, which links households to different income and wealth states temporally, can lead to mechanisms being missed.

Population aging and rising inequality are placing downward pressure on interest rates in many advanced economies. This section illustrates some of the subtleties involved in examining the distributional impact of these changes. Future work should analyze the impact of the monetary and fiscal policy responses to these effects. Further, while this section focused on the long-term distributional impact, it would also be fruitful to analyze the effects along transition paths resulting from changes in safe interest rates.

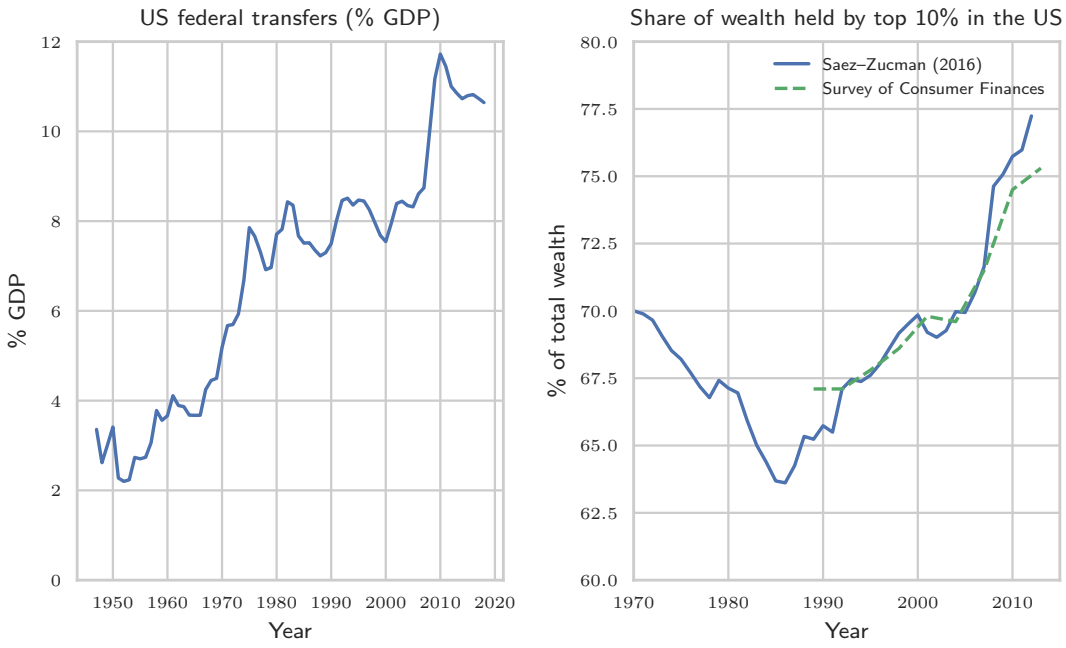
Part III

CHAPTER 4

REDISTRIBUTION FROM THE CRADLE TO THE GRAVE

FIGURE 4.1: INCREASING TRANSFER SPENDING AND WEALTH INEQUALITY IN THE U.S.

Data: BEA, SCF, Saez and Zucman (2016)



4.1 INTRODUCTION

Many countries have seen large changes to fiscal policy over the past 50 years, with fluctuations in tax rates and increases in social insurance expenditure. Often – notably in the United States – these changes were followed by a rise in wealth

inequality. There is an ongoing debate about whether these issues are linked. Key to this question is understanding how household savings respond to fiscal changes, and what this implies for the distribution of wealth.

In this chapter, I provide a quantitative model of the dynamic effects of fiscal reform on household savings, the wealth distribution, and macroeconomic aggregates.

Public transfers can crowd out private savings: transfers serve as external insurance to households against fluctuations in their income, alleviating the motive to accumulate wealth as self-insurance. The extent of this crowding out varies across households with different wealth and income, and at different points in their lifecycle. As a result, a change to social insurance has a heterogeneous impact on the distribution of household savings. Over time, this leads to changes in the distribution of household wealth.

A simple example illustrates these dynamic effects on the wealth distribution. Suppose a policymaker expands a program of (possibly means-tested) transfers to senior citizens, such as the Supplemental Social Insurance program existing in the United States. (We ignore for now the source of funds for these extra transfers.) The short-run effect of this reform is likely to be a small reduction in wealth inequality, as poor seniors receive more wealth. However, over the long run, working households incorporate this transfer schedule into their consumption–savings decision. Poorer workers prefer to consume more during working life and rely on these extra transfers for retirement consumption, thus accumulating less wealth over their lifecycle. By contrast, these transfers do not substantially affect the savings of the rich, being small relative to their assets. Thus, this reform – designed to increase social insurance – will lead to higher wealth inequality over

the long-run, an idea from Hubbard et al. (1995). Understanding how – and how fast – we get from a slight decrease in inequality to a longer-term increase is one of the key questions of this paper, which I address by means of a quantitative model of the dynamics of the wealth distribution.

In light of increasing wealth inequality, it is important to be able to quantify the distributional impact of different channels. Firstly, economists are debating the causes of the rise in inequality – especially in the United States – weighing the contribution of structural changes (skills-biased technological change leading to rising income inequality) and fiscal changes (with a focus on declining taxes on the rich since the 1960s).[1] In this paper, I find that it takes around 25 years for the full effect of social insurance on wealth inequality to materialize, suggesting that the contribution of social insurance to wealth inequality may have been partially missed due to the lag. Indeed, Pham-Dao (2016) finds that differences in social insurance account for 70% of the variation in wealth inequality across Eurozone countries.

Rising wealth inequality highlights a second question. What – if anything – is the appropriate policy response? Redistributive transfers may reduce wealth inequality in the short term but are likely to *increase* it over the long-run through the channel discussed above. Further, social insurance is generally welfare improving at the margin, so that the accompanying wealth inequality can be seen as a side effect to ignore. Indeed, despite popular backlash against wealth inequality (a natural result of political economy considerations), the policymaker would do better to focus on welfare measures, permanent income inequality, or consump-

[1] Recent discussions include Kaymak and Poschke (2016), Hubmer et al. (2016), and Heathcote et al. (2017).

tion inequality. However, regardless of the policymaker's objective function, a tool for predicting the distributional impact of fiscal policy remains useful.

To this effect, I build a macroeconomic framework with which to analyze these distributional effects, nesting the standard incomplete markets model and canonical lifecycle mechanisms. Using this model, I make two main contributions in this chapter.

First, I quantify the cross-sectional reaction of household savings to social security reform in the model. The response is highly heterogeneous, with 45–65 year-olds in lower wealth groups dissaving the most. These patterns are in line with a growing body of empirical evidence on how public pensions crowd out private wealth.[2] Crowding out of savings is concentrated strongest at the bottom of the wealth distribution. As a result, the change in aggregate wealth is small, so that the quantitative effect on the equilibrium interest rate is negligible. By addressing this question with a dynamic general equilibrium model, I thus find additional support for a large applied and empirical literature that analyzes the effects of social insurance on consumption and savings from a structural partial equilibrium standpoint.[3]

Second, I compute the dynamics of the wealth distribution resulting from the change in household savings. I find (in line with the theoretical results of Hubbard et al. (1995)) that more social insurance leads to higher wealth inequality. In fact,

[2]Bottazzi et al. (2006), Attanasio and Rohwedder (2003), Aguila (2011), Feng et al. (2011), and Lachowska and Myck (2018) look at various pension reforms in the 1990s in Italy, the United Kingdom, Mexico, China and Poland respectively, finding evidence for relatively strong crowding-out effects using a difference-in-difference design. By contrast Chetty et al. (2014) use a regression discontinuity approach around a reform based in Denmark and find much weaker effects. I discuss this literature in more detail in Section 4.3.4.

[3]See, for example, Deaton et al. (2002), Storesletten et al. (2004), Heathcote et al. (2009, 2014, 2017), and Huggett et al. (2011).

these dynamics are slow: it takes roughly 25 years for the full distributional effect to materialize. The reason for this is twofold. First, it takes time for changes in savings to cumulate to changes in the stock of wealth. Second, there are temporary cohort effects if the extra transfers are financed contemporaneously with taxes, as early generations receive transfers which they have not fully paid for.

Accordingly, I analyze how the reactions of savings and wealth vary for different kinds of transfer programs. Making transfers more redistributive strengthens the distributional effects: more progressive subsidies lead to even *higher* long-run wealth inequality, by increasingly crowding-out specifically the savings of poorer households. I also consider means-tested transfer programs aimed at younger workers. These insure households against income shocks rather than lifecycle variation in earnings. Subsidies to workers have a weaker effect on aggregates than transfers to seniors, as households find it easier to save against future taxation than to borrow against future transfers. These transfers also tend to increase wealth inequality, by reducing the savings of households with low earnings.[4] Moreover, better insured against income risk, households are willing to take on riskier portfolios, further increasing wealth inequality.

This paper uses a unified model of heterogeneity in age, income and wealth based on the framework of Chapter 1. The model is parsimonious by design, the union of a canonical lifecycle model with a standard incomplete markets model of heterogeneity in income and wealth. Lifecycle motives come from uncertain lifespans – determined from actuarial data on survival rates – and a hump-shaped profile of average income over the lifecycle, obtained from the data in the Sur-

[4]Income is positively correlated with wealth, especially early in the lifecycle.

vey of Consumer Finances. Households face uninsured idiosyncratic shocks to income at each age.[5] The model yields empirically realistic lifecycle patterns and distributions of income and wealth despite its parsimony, as we saw in Section 1.3.

Literature

This paper contributes to several strands of the literature.

There exists a substantial empirical and applied literature which analyzes the effect of social insurance on household consumption and savings, using partial equilibrium models of lifecycle behavior and incomplete markets.[6] This paper builds on that literature by analyzing the resulting dynamic distributional effects and the accompanying general equilibrium changes. I find that the reaction of the interest rate is quantitatively small, providing support for the partial equilibrium approach taken by these studies.

Several papers investigate the effects of fiscal policy on the wealth distribution. Kaymak and Poschke (2016), Hubmer et al. (2016), and Heathcote et al. (2017) focus on the relationship of top income tax rates to tail wealth inequality. By contrast, this paper emphasizes the distributional effects of social insurance, first argued by Hubbard et al. (1995) to increase long-run wealth inequality. Using a partial equilibrium structural model, Pham-Dao (2016) finds that social insurance can explain 70% of the variation in wealth inequality across the Eurozone.

Finally, I take the approach in this paper of building a parsimonious model nesting canonical mechanisms of lifecycle behavior and incomplete markets. The

[5]The model does not feature aggregate risk for technical reasons, although Ahn et al. (2018) provide promising results in that direction.

[6]See for example Gourinchas and Parker (2002), Storesletten et al. (2004), Kaplan (2012), Deaton and Paxson (1994), Cocco et al. (2005), Ziliak (2003), Deaton et al. (2002), Heathcote et al. (2009, 2014, 2017), Huggett et al. (2011), and Hubbard et al. (1995).

idea is to use the extra computational power afforded by the new solution method to enlarge the household state space, and to see how well a standard model matches the data in such a setting. As a result, there are mechanisms not explicitly captured in the model that may be relevant to this question. First is health risk, a form of uncertainty displaying cross-sectional and lifecycle patterns that has been argued to be important for household savings (Gruber and Yelowitz 1999; Maynard and Qiu 2009; De Nardi et al. 2016). Health expenditure could be introduced to the model as a form of negative income risk to seniors. Faced with this risk, households have a larger requirement for insurance. Public transfers may thus crowd out their savings less strongly, so that the elasticities computed in this paper can be seen as an upper bound.

Another important determinant of household savings behavior is the bequest motive.[7] These are important for matching certain empirical features of current household savings. However, when analyzing *changes* in saving (i.e. crowding out), the effects of bequests seem to cancel out. I thus omit bequests from the version of the model used in this section, but Chapter 1 explains how to reintroduce them.

The remainder of this chapter is laid out as follows. Section 4.2 explains the model. Section 4.3 sets out the distributional effects of a variety of reforms to social insurance, and Section 4.4 concludes.

4.2 model

In order to examine the impact of social insurance on household savings, we need a tractable model incorporating the heterogeneity in savings coming from

[7]See, for example, Castaneda et al. (2003), Benhabib et al. (2011), and Straub (2017).

the lifecycle and from variations in income and wealth. To that end, I build a heterogeneous agent overlapping generations model, running in continuous time, which is a special case of the framework set out in Section 1.2. Households face idiosyncratic risk to their income and to their capital returns. Their labor income fluctuates around a hump-shaped profile of income against age, taken from the data, which leads to lifecycle savings behavior.

Relative to Section 1.2, I make one major simplifying modeling choice. I disable bequests, setting the bequest motive χ to zero. Numerical experiments indicate that bequests have a negligible quantitative effect on the main results of this chapter, while imposing a significant computational burden. As a result, I solve for an equilibrium without intergenerational transmission (as defined in Section 1.2.4). Details of the model can be found in Section 1.2, and I set out the key equations below for the sake of clarity.

Without a bequest motive, households invest maximally into the annuity. The budget constraint for a given household is thus given by

$$\mathrm{d}w_t = (y_t z_t + \Theta(t, w_t, z_t) + r_t^b w_t + (r^k - r_t^b)k_t - c_t + \pi_t w_t)\,\mathrm{d}t + \sigma k_t\,\mathrm{d}B_t. \quad (4.1)$$

The Hamilton–Jacobi–Bellman equation for the value of a household at time τ with age t, financial wealth w, and productivity z reduces to

$$\begin{aligned}(\rho + \lambda_t)V(\tau, t, w, z) = \max_{c,k}\, & u(c) + \partial_t V + \partial_\tau V \\ & + \partial_w V\big[y_t z + \Theta(t, w, z) + r_t^b w + (r^k - r_t^b)k - c + \pi_t w\big] \\ & + \partial_{ww} V \frac{\sigma^2}{2} k^2 \\ & + \partial_z V \mu^z(z) + \eta \int_{-\infty}^{\infty} (V(t, w, z') - V(t, w, z))\phi^z(z')\,\mathrm{d}z'.\end{aligned} \quad (4.2)$$

The dynamics of the joint distribution of household age, wealth and productivity $g(\tau, t, w, z)$ at time τ are given by the *Kolmogorov Forward equation*

$$\begin{aligned}\partial_t g(\tau,t,w,z) =& \delta(t)g_0(w,z) - \lambda_t g - \partial_w[sg] + \partial_{ww}\left[\frac{(\sigma k)^2}{2}g\right] \\ &- \partial_z[\mu^z g] + \eta\int_{-\infty}^{\infty}\big(g(\tau,t,w,z') - g(\tau,t,w,z)\big)\phi^z(z')\,\mathrm{d}z',\end{aligned} \tag{4.3}$$

where $g_0(w,z)$ is the distribution of wealth and productivity for newly formed households, and δ is a Dirac delta function.

Together with the bond market clearing condition,

$$\int\int\int (w - k(\tau,t,w,z))g(\tau,t,w,z)\,\mathrm{d}t\,\mathrm{d}w\,\mathrm{d}z = 0 \quad \text{for all } \tau, \tag{4.4}$$

these equations characterize the model, specifying an equilibrium without intergenerational transmission. The solution methods of Chapter 1 can be directly applied to this setup.

In Section 1.3 I set out the calibration for this model, and evaluated its ability to match moments in the data relating to the lifecycle and to the distributions of income and wealth. Fiscal policy consists of a flat transfer to seniors (households 65 or older) of \$10,000 per year, financed using a progressive income tax. I show that the parameterization reproduces salient lifecycle and distributional features of the data, despite its relative parsimony. In numerical experiments, I find that what matters for the main results of this chapter are the reactions to *changes* to fiscal policy rather than the initial setup itself.

4.3 Distributional effects of social insurance

In this section, I examine the impact on household savings and on the wealth distribution of reforms to social insurance. I focus on a variety of retirement

transfer programs, and additionally consider income-tested subsidies to workers at the end of the section.

I look at how this change in transfers affects consumption and saving across age, wealth and income groups in the model, and how these effects vary over time. Responses are very heterogeneous: some households essentially do not react, while the savings of others change with large elasticities. The magnitude and distribution of these reactions align with the empirical literature estimating reactions of savings to a range of fiscal policy changes.

The cross-sectional changes in savings lead to large and surprising effects on wealth inequality. In particular, more generous social security leads to increased wealth inequality, especially among seniors.

The first experiment examines a marginal increase in transfers to seniors, financed by an income tax hike. This generalizes the textbook exercise in which pay-as-you-go social security is introduced to an overlapping generations economy. There, the two takeaways are that (a) there are winners and losers (the current old get a free ride) and (b) this policy can be Pareto improving by remedying dynamic inefficiency. Here, the answer to (a) becomes more nuanced: the current old still get a free ride, but many current workers also become better off by the provision of this insurance (despite higher taxes).

The full effect of the mechanism takes the length of one career to be felt. When the policy is announced, workers rethink their consumption and retirement savings decision. Faced with more safe retirement income, they want to save less and consume more now. The extent of this change depends on their wealth and income. Asset-rich households plan to use their savings to consume a lot at retirement; transfers comprise a small part of their retirement income. As a

result, the change in transfer does not significantly impact their savings decision.

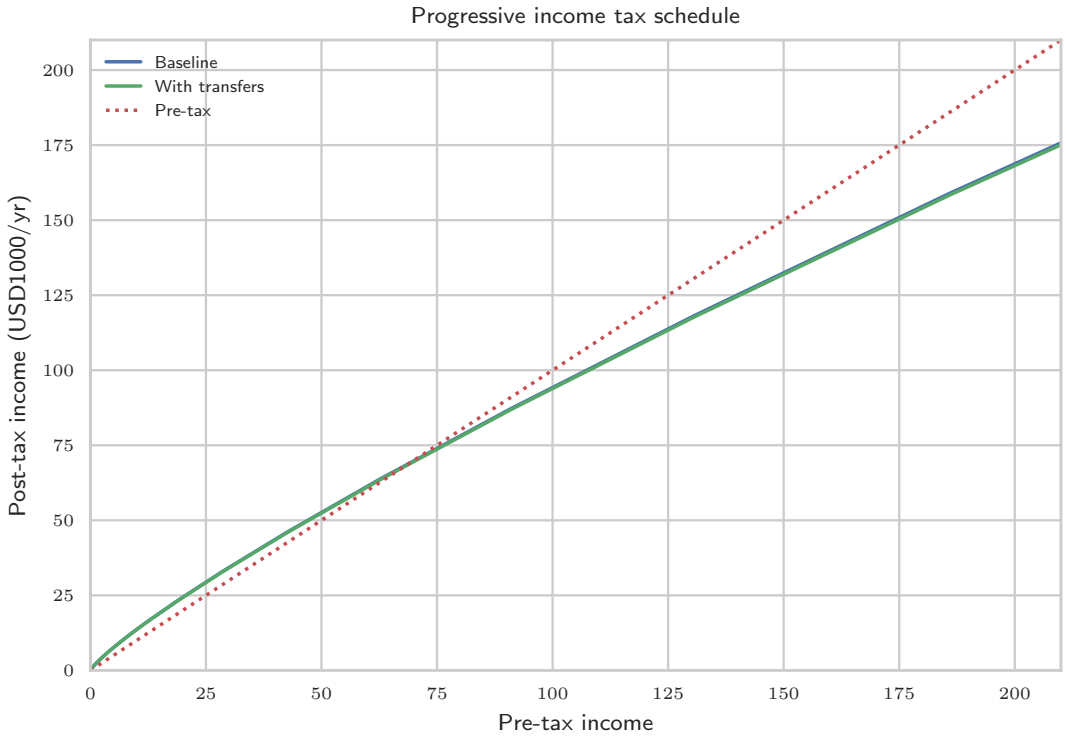
Less well-off workers, however, find themselves with more permanent income coming later in life. They prefer to sacrifice saving for more consumption today, accumulating less wealth. The poorest workers in the model are hand-to-mouth consumers. They already are not and cannot not save.

As a result, we see a highly nonlinear reaction of savings across the wealth cross-section. Households in the top end of the wealth distribution are not significantly affected by the policy change, while those in the middle and lower end of the distribution accumulate less wealth. As a result, wealth inequality increases. We shall see later that this increase is driven by changes in the lower half of the wealth distribution – tail inequality is left mostly unchanged.

There are further effects to consider. When the pension increase is paid for by higher taxes, these affect wealth groups differently. Taxes aside, other mechanisms also link these transfers to inequality, such as a portfolio choice channel and interaction with buffer stock saving. Finally, those who were already alive at the time of the policy announcement will find themselves with different wealth patterns and elasticities as later generations, leading to medium-term dynamic effects. I discuss all of these lower down.

Later in this discussion, I compare the reaction to this flat social security increase to the case of an asset-tested program. I find that progressivity of the reform greatly amplifies the distributional effects, with important policy implications. I also run similar experiments for transfer programs targeting workers. There, the reactions of aggregates are less strong, as workers anticipate future tax liability. In Section 4.A.1, I detail further policy experiments to demonstrate the robustness of the results in this section.

FIGURE 4.2: INCOME TAX PROFILES: BASELINE AND RAISED FOR EXTRA TRANSFERS



4.3.1 *Setup*

In the initial equilibrium, social security consists of a universal lump-sum transfer of \$10,000/year to all agents over 65. This transfer is paid for by the progressive income tax described in Section 1.3.2. In the main experiment, fiscal policy changes as follows: the social security transfers increase by 10%, and the level of the income tax θ_0 increases in order to keep fiscal balance, keeping the progressivity parameter θ_1 fixed. The income tax schedule in the baseline and in the experiment are plotted in Figure 4.2. I analyze further pension reform exercises in appendix section 4.A.1.

The fiscal reform is a surprise, announced and implemented at calendar date $\tau = 0$. In order to tease out the effects of the fiscal change, I decompose the reaction into three parts:

1. the partial equilibrium with higher transfers but *without* the associated tax hike (in which the extra transfers arrive for free, like manna from heaven);
2. the partial equilibrium with higher transfers and the higher income tax;
3. the general equilibrium response to both higher transfers and higher taxes, when r^b adjusts.

The incremental effect of part (3) is identical to the discussion of Section 3.4, with the equilibrium rate here increasing slightly. For this reason, I focus on (1) and (2) above.

The empirical literature focuses on the coefficient of substitution of private savings for expected pension wealth. Pension wealth in the model is defined as the expected present value of pension benefits,

$$PW(t,w,z) = \mathbb{E}_t\left[\int_t^\infty \Theta(t',w_{t'},z_{t'})1_{t'\geq 65}e^{-\int_{t'}^{\xi}(r_s^b+\lambda_s)\,\mathrm{d}s}\,\mathrm{d}t' \mid w_t = w, z_t = z\right]. \tag{4.5}$$

Here $\Theta(t',w,z)1_{t'\geq 65}$ are the transfers received after age 65. Section 4.A.2 explains how to calculate this expectation efficiently using a Feynman–Kac equation.

4.3.2 *Comparative statics*

I begin with looking at the long-term effects of increasing social security, by comparing the high-social-security and low-social-security steady states. Since the fiscal reform hits agents of all ages, it takes a long time to converge to the

Table 4.1: Change of aggregates over pension wealth change, with and without corresponding tax increase

Variable	(1) No Tax	(2) Tax (Partial Eq.)	(3) Tax (General Eq.)
Wealth	-54%	-67%	-69%
Savings (*flow*)	-2.7%	-3.1%	-3.2%
Cons.	2.5%	-0.5%	-0.8%
Wealth Gini*	0.0279	0.0281	0.0263

* Elasticity

new steady state after the reform is announced, on the order of a lifetime. We shall examine the transition dynamics in the next section.

Table 4.1 lists the reaction of model aggregates and inequality to the increase in social security, computed as the difference from the baseline aggregates divided by the change in pension wealth.

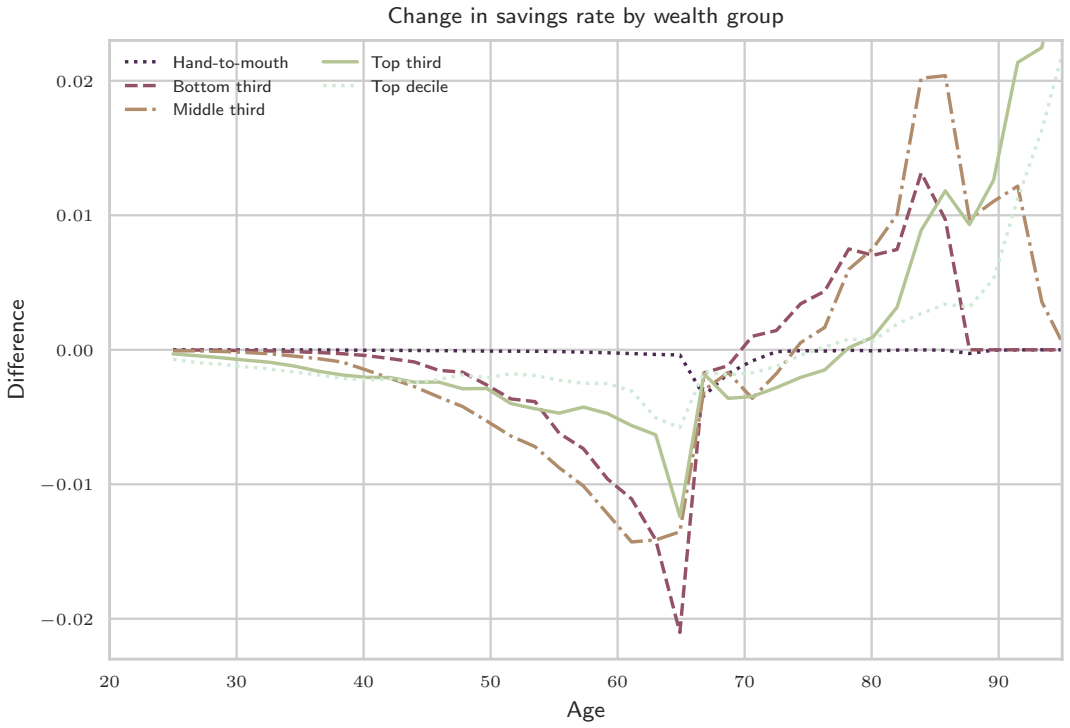
The first column lists elasticities in the case of higher social security transfers without matching tax increase. There is a large amount of crowding out: an extra dollar of pension wealth reduces household financial wealth by 54 cents. This is achieved by a drop in savings, smoothed out over working life. As a result, consumption increases. The resulting changes in the wealth distribution will be discussed shortly.

The second column has the transfer increase together with the tax increase, but in partial equilibrium (where the interest rate r^b is fixed at the baseline level and the bond market does not clear). Relative to the No Tax scenario (1), the wealth effect of the income tax hike leads to stronger negative reactions of aggregate savings and wealth. These feed through to consumption, which also falls.

Finally, the third column outlines the full general equilibrium response with

Figure 4.3: Changes in savings rate in response to social security increase

Note: Plots savings divided by total disposable income over the lifecycle in response to a 10% increase in social security, for different wealth groups. (Averaged over income groups, *without* accompanying tax hike.)



transfers and tax hike. The bond market clearing interest rate r^b increases from 2% to 2.04%; an elasticity of 0.20 to the transfer increase. Relative to the second column, savings and wealth fall further, along with consumption. The wealth distribution becomes less unequal, as we shall see in Section 3.4.

We can tease out a fuller picture by looking at cross sections. Figure 4.3 plots the change in savings rate for each age group, defined as savings over total post-fiscal income. Agents save less during working life, especially so right before retirement where they smooth consumption using their savings, and anticipate

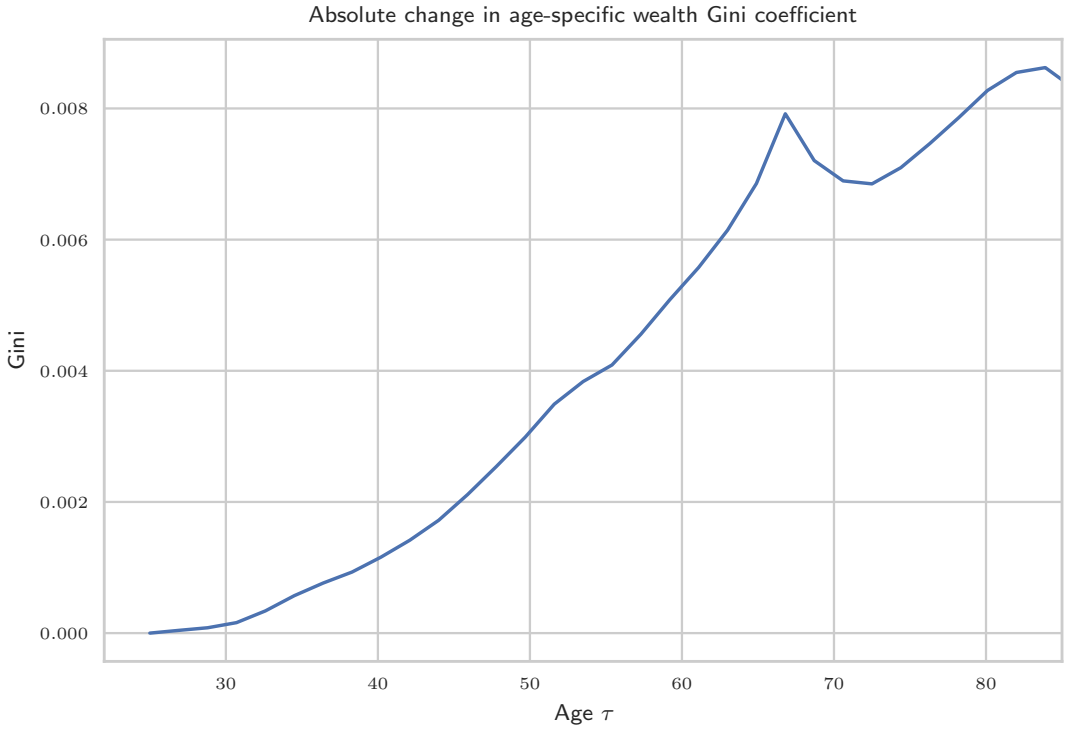
the extra retirement transfers. Those in the middle and bottom third of the wealth distribution dissave the most. For them, these extra transfers are an important component of anticipated retirement income, replacing their need to self-insure. By contrast, the hand-to-mouth and the top 10% do not change their savings substantially, either because they cannot or because for them the extra social security is insignificant.

At age 65, the savings rate of all wealth groups is a few basis points lower than in the baseline. As they age, the picture is almost a mirror image of the graph before retirement, with poorer agents increasing their savings rate the most. This is for a reason. Agents who reduced their retirement savings during working life now enter retirement with less financial wealth. As a result, their capital income in retirement is necessarily lower. This pushes up their savings rate for a given level of savings. Although the level of savings shifts down (and consumption shifts up) as they consume the extra transfers, the impact on income dominates.

The consumption–savings decision for retirees is simpler. Social security does not vary with age; the only source of lifecycle variation for seniors is mortality risk. As a result, the transfers are less insurative than they would be during the more volatile period of working life. Retirees find themselves with more permanent income which is immediately available. Consumption and savings both increase; especially so for retirees in lower wealth groups.

Working life is when savings behavior has the most lasting impact on wealth accumulation. Poor workers dissave more than rich workers, due to the differential rates at which the public transfers crowd out private savings. As a result, wealth inequality increases. These channels are illustrated in Figure 4.4, which shows that wealth inequality by age group increases across the board, with effects

Figure 4.4: Long-run change in wealth inequality by age group after a social security increase



cumulating into larger changes later in the lifecycle.

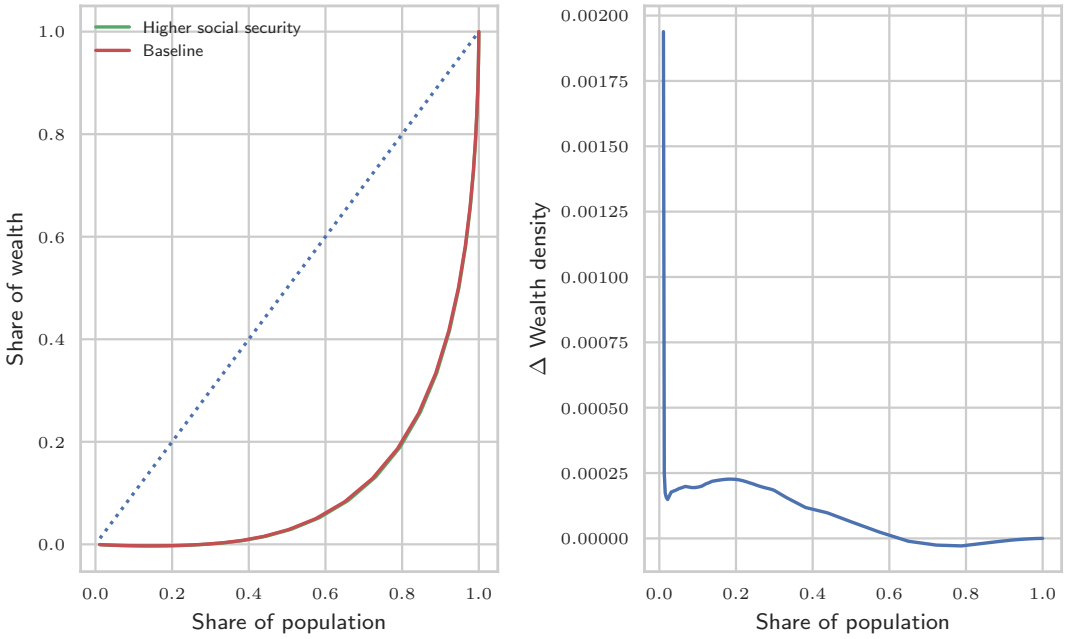
Another channel for the increase in wealth inequality is portfolio choice. With more safe pension income, household consumption becomes less sensitive to fluctuations in net worth. This serves to reduce the effective level of risk aversion of agents and thus to encourage investment in the risky asset, as we shall discuss in detail in equation (3.4). Since k carries *idiosyncratic* risk, this will lead to a further increase in wealth inequality.

These channels are in fact two sides of the same coin, underlining the insurative role for households played by public transfer policies.

Which parts of the wealth distribution are affected? If crowding out dispro-

FIGURE 4.5: LONG-RUN IMPACT OF SOCIAL SECURITY INCREASE ON THE MARGINAL DISTRIBUTION OF WEALTH

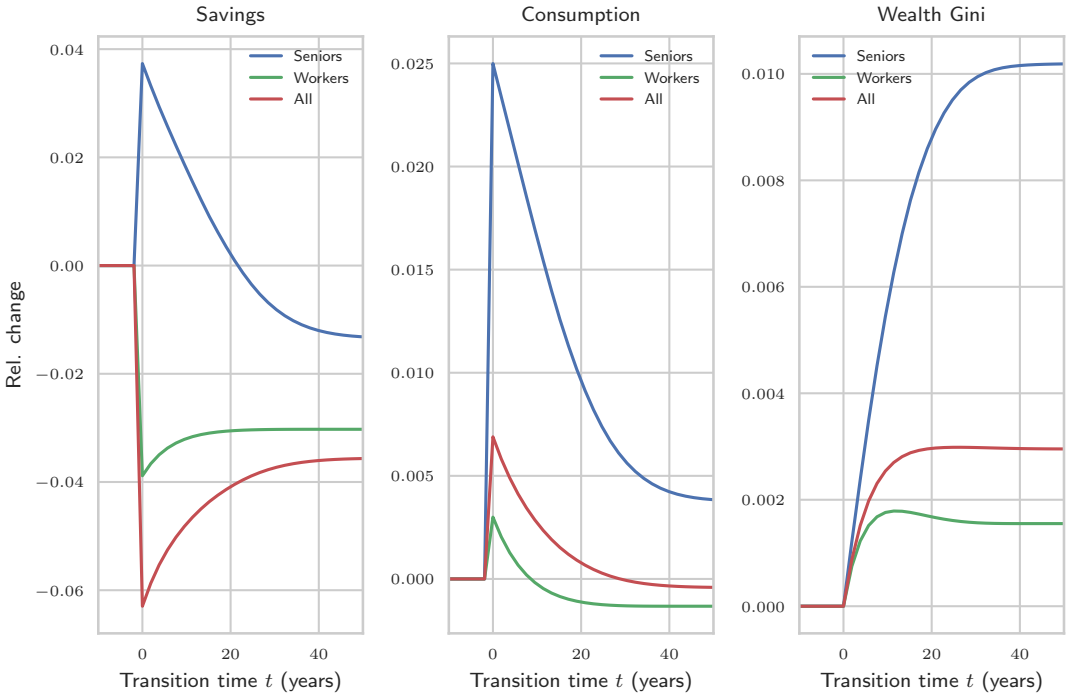
Left panel: Lorenz curves for the two steady states. Right panel: Change in wealth density at different percentiles.



portionately reduces the savings of poorer households, we would expect most of the change in the wealth distribution to be concentrated at the bottom. Figure 4.5 illustrates this idea. The Lorenz curve plots the share of total wealth (vertical axis) held by those in the bottom $x\%$ of the wealth population (x-axis). Although the Lorenz curves for the two distributions seem indistinguishable, the right panel illustrates that more mass is added at the bottom half of the distribution, with a spike in the number of agents against the borrowing constraint.

Figure 4.6: Dynamic response of savings, consumption, and wealth inequality to social security increase

Note: Effect of a 10% increase in social security. Workers and retirees are defined as households with age below and above age 65, respectively.



4.3.3 *Dynamic effects*

Figure 4.6 shows how aggregates and wealth inequality evolve through the transition from the surprise announcement at time 0 of the increase in transfers and taxes. The first thing to note is that convergence is slow, on the order of 50 years.

As in the textbook analysis of pay and you go social security in an overlapping generations model, the winners are the current old, who greatly increase both consumption and savings when they receive a windfall which they have not paid

for. Further, current workers generally increase consumption too (and reduce savings) as all but the youngest workers will be receiving pension income that future generations will have to pay for.

Over time, workers will start to save more: they expect to spend more time paying the higher income tax and so the transitory bump to their permanent income tapers off. Similarly, after around $\tau = 25$ years, many of the current retirees will be agents who have paid the higher income tax for a significant portion of their working life. As a result, they enter retirement with less financial wealth and thus less capital income than earlier retirees. Their savings gradually fall below those of retirees in the initial steady state. Their consumption remains higher thanks to their higher pension wealth.

Due to these effects on savings, wealth inequality increases gradually across all groups. It reaches an apex around $\tau = 25$, after most of the windfall-receiving generation has died and enough time has passed for the new retirees to have had their savings in working life substantially affected by the reform.

4.3.4 *Empirical estimates*

There is a large empirical literature estimating how much public pensions crowd out private savings. [8] These studies estimate the degree of substitution between the household savings rate and to pension wealth (expressed in units of current income). They find this coefficient to lie broadly between -0.2 and -0.7. This is a measure of substitution of a flow variable for a stock variable. The estimates are

[8] For example, Attanasio and Brugiavini (2003), Attanasio and Rohwedder (2003), Aguila (2011), Feng et al. (2011), and Lachowska and Myck (2018) use a difference-in-difference design to estimate the crowding out effects of pension reform exercises in Italy, the UK, Mexico, rural China, and Poland, respectively.

surprisingly large: an extra unit of pension wealth leads savings *each year* to be 0.2 to 0.7 units lower. In other words, extra pension wealth will crowd out the stock household wealth more than one to one within five years.

Bottazzi et al. (2006) address this issue by estimating the reaction of the stock of household wealth to changes in pension wealth, studying three pension reforms in Italy in the 1990s. They find a degree of substitutability around 65% overall, with a higher coefficient (around 80%) for agents well-informed about their pensions.

By contrast, Chetty et al. (2014) uses a regression discontinuity design to analyze the impact of a pension reform in Denmark in 1998, in which firms became mandated to make automatic contributions to workers' retirement accounts for earnings above a threshold. They find no significant relationship between private saving and public pension wealth. One suggested explanation for the difference (Lachowska and Myck 2018) is that the reforms studied by Chetty et al. (2014) and Feng et al. (2011) *increase* pension wealth, while the other studies look at reforms that reduced it. This theory is supported by the increasing evidence of a large asymmetry in micro-level propensities to save and consume – reactions to negative shocks can be an order of magnitude larger than to positive (Bunn et al. 2018). Further, the discontinuity design relies on the savings behavior of agents around the threshold (earnings of $5,300), while other studies use a broader design. Finally, Chetty et al. (2014) finds a negligible impact among most agents, but a large reaction conditional on a reaction of savings. It may be that the reform in Denmark, which entailed automatic pension contributions from employers, went unnoticed by many households.

In the model, substitutability is around 0.7, on the higher end of available es-

timates. This measure is closer also to estimates for well-informed agents derived by Bottazzi et al. (2006). This makes sense, given the well-known inattention to pension arrangements in the data (Thaler and Benartzi 2004; Chetty 2015). Indeed, the extent to which the public were aware of pension reforms may explain the large variation in substitutability across different countries. In that sense, the crowding-out measures of this paper can be seen as an upper bound.

Another useful source of estimates is the literature on the impact on savings of reforms to public health insurance. Maynard and Qiu (2009), De Nardi et al. (2016), and Gruber and Yelowitz (1999) study reforms to Medicaid, the asset-tested medical insurance program in the U.S. They find a range of elasticities of private savings to changes in the transfer program between 0 and -0.15. Additionally, Abbott et al. (2013) find evidence that education subsidies in the United States strongly crowd out parental transfers.

These studies do not present sufficient cross-sectional evidence to fully validate the model's predictions for savings across wealth groups. However, distributional evidence on marginal propensities to save and consume out of windfalls[9] provide support for the findings of the model, identifying large reactions in the bottom half of the wealth distribution that decline with net worth and with age.

4.3.5 *Progressivity of social security*

The experiment above analyzed flat reform to public pensions, increasing transfers to all senior-aged agents equally. In this subsection I look at the effect of a *progressive* reform to social security, by introducing a means-tested subsidy

[9]For example, Misra and Surico (2014), Kaplan and Violante (2014), and Kaplan et al. (2014) study fiscal windfalls, Fagereng et al. (2016) look at lottery winnings, and Bunn et al. (2018) use a survey.

for the poorest retirees. Many countries have such programs in place, such as the Supplemental Security Income (SSI) in the United States, which pays up to $1,100 per month to a couple with net worth below $3,000. In order to retain comparability with the exercise above, I look at a marginal version of this program which adds a small means-tested component to an existing social security setup. I find that progressivity amplifies all the mechanisms set out in the last section.

I study a reform that introduces an extra transfer schedule $\tilde{\Theta}$ which adds $950 per year to all agents plus a gradual subsidy for households with net worth below $1,000,

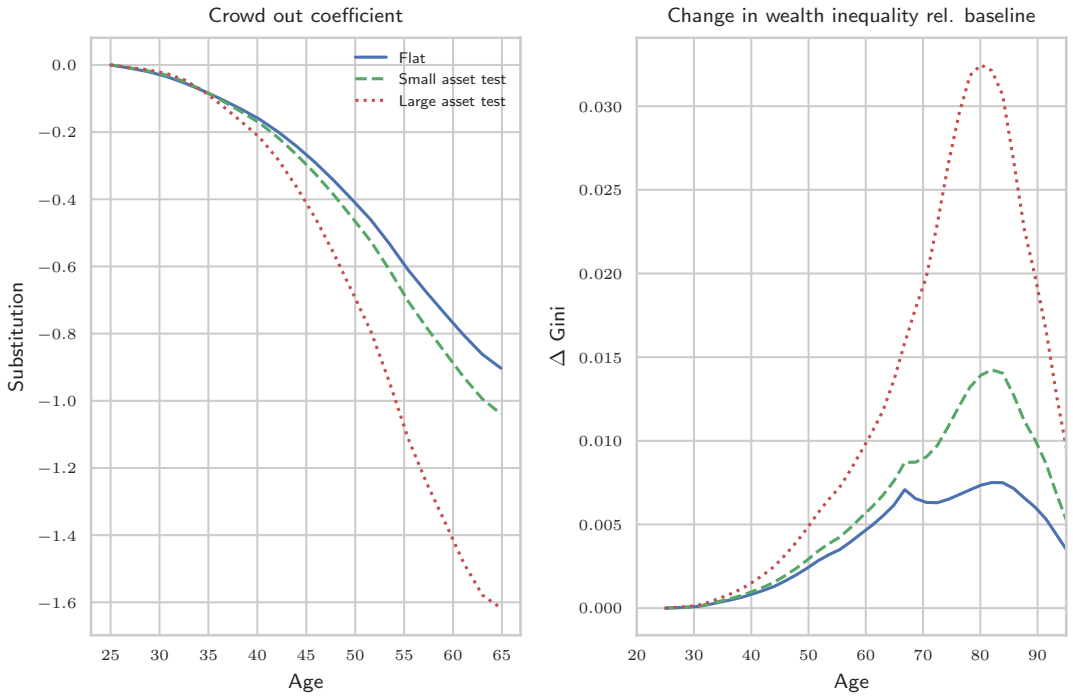
$$\tilde{\Theta}(w) = \begin{cases} 950 + 0.4(1000 - w) & \text{if } w < 1000; \\ 950 & \text{otherwise.} \end{cases}.$$

This way, the reform remains broadly similar to the setup analyzed above, with the addition of a small wealth tested top-up. For comparison's sake, I will also discuss a reform with a larger means test, where

$$\tilde{\Theta}^{\text{big}}(w) = \begin{cases} 800 + 1.5(1000 - w) & \text{if } w < 1000; \\ 800 & \text{otherwise.} \end{cases}.$$

Progressivity amplifies the asymmetry in savings reaction through cross-section, illustrated in Figure 4.7. As before, private savings are crowded out by these extra transfers. The transfers are also seen as a particularly effective form of insurance for the poor households targeted by the policy. As a result, there is a high degree of substitutability of their savings for transfers. This discouragement is so strong that over the lifecycle the larger means-tested program crowds out private wealth more than one to one. Further, asset-tested transfers carry an implicit wealth

FIGURE 4.7: EFFECT OF MEANS-TESTING SOCIAL SECURITY TRANSFERS ON SAVINGS AND THE WEALTH DISTRIBUTION, OVER THE LIFECYCLE



tax which strongly discourages savings for households near or below the eligibility threshold of the program. On the other hand, better-off households with no hope of eligibility will substitute relatively less.

As a result, the distributional impact of this progressive reform is also much larger, with a three- or six-fold higher increase of the age-specific Gini coefficient for wealth than in the baseline lump-sum case.

Reactions of aggregates are also amplified, but to a smaller extent. Although the savings of poorer households fall drastically in relative terms, these account for a smaller fraction of aggregate savings and consumption. Indeed, the reaction of savings is only slightly increased, but this feeds through to a large cumulative

effect on aggregate wealth.

This example highlights the sensitivity of distributional variables to fiscal reform, even as aggregates display a more mild reaction. It is crucial, therefore, to evaluate the (often explicitly distributional) objectives of fiscal policy in a setting that captures the nonlinear and heterogeneous reaction of households.

4.3.6 *Worker subsidies*

To what extent do the effects above carry over to subsidies to *workers* instead of retirees? Rather than providing insurance against lifecycle variation in income, worker subsidies mainly affect the precautionary savings motive. I highlight two scenarios in this section, and consider others in Section 4.A.1. The programs are selected to cost the same to the government as the social security increase, so that the accompanying tax increase is identical.

Means-tested programs

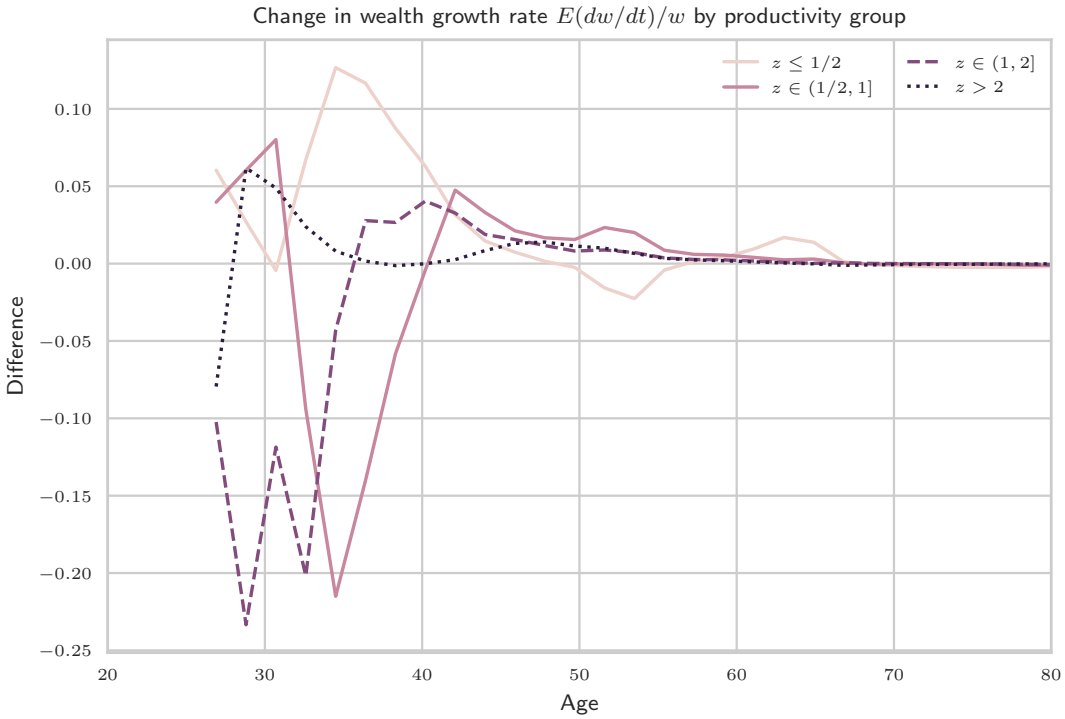
The first scenario is a productivity-tested transfer program. Since there is predictable lifecycle variation in earnings, purely income-tested benefits redistribute across the age groups as well as income groups. A productivity test allows us to isolate effects on the precautionary savings motive. These are equivalent to having benefits vary with the ratio of a household's income to the median income for households of that age.

The two means-tested policies I examine take the form

$$\text{Small productivity test} = 1_{\tau \leq 65} \max\left\{0, 1500(1.0 - z)\right\},$$

$$\text{Large productivity test} = 1_{\tau \leq 65} \max\left\{0, 700(0.5 - z)\right\},$$

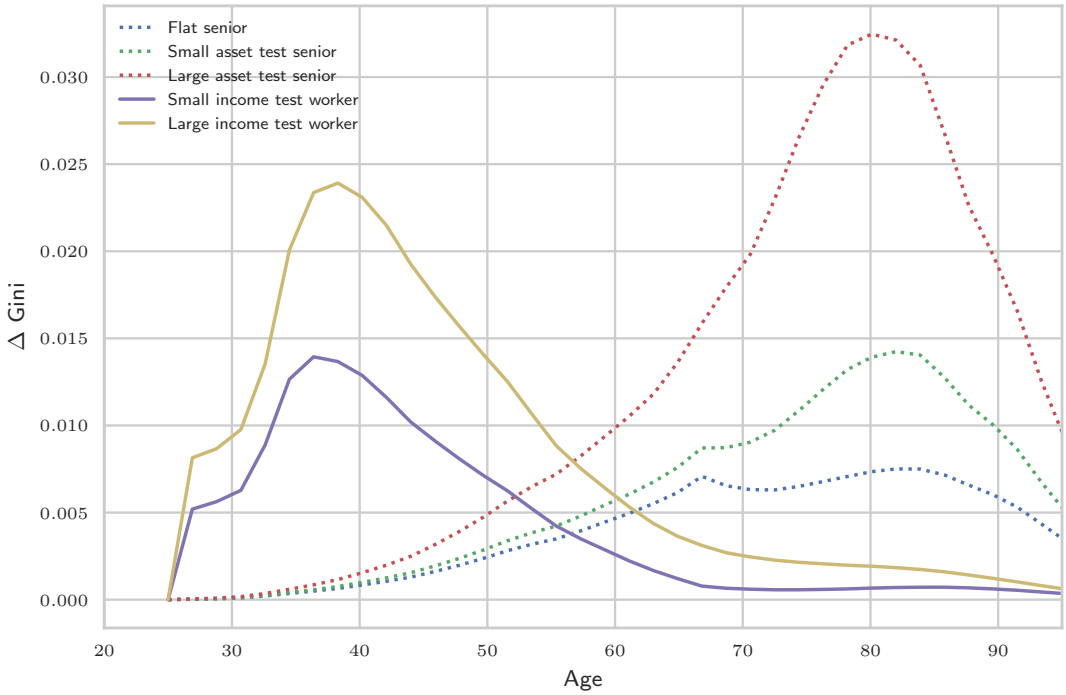
FIGURE 4.8: EFFECT OF MEANS-TESTED TRANSFERS TO WORKERS ON LIFE-CYCLE WEALTH ACCUMULATION



designed to cost the same to the government as the pension reform exercises considered above ($1,000p.p.).

Figure 4.8 plots the long-run effect of the small productivity-tested transfers to workers on wealth accumulation through the lifecycle, plotted for different productivity groups. The extra transfers allow the least productive agents ($z \leq 1/2$) to accumulate more wealth. By contrast, for households close to the eligibility margin of 1 (i.e. those with $z \in (1/2, 2)$), these transfers are viewed as an effective form of insurance against income shocks. This strongly crowds out their precautionary savings motives and causes them to dissave. The most productive agents ($z \geq 2$), unlikely to become eligible for the program, do not substantially

Figure 4.9: Effect of various policies on lifecycle wealth inequality



change their savings behavior. Since wealth and income are correlated (especially in early life), we thus see a similar differential crowding out effect at play which will lead to more wealth inequality.

Note that for all groups, the effect on savings becomes small after age 40. This finding is in line with Gourinchas and Parker (2002), who find that empirically households save mostly for precautionary reasons until age 40 and then switch to saving for retirement. Similar behavior is at play in the model, with households facing relatively large income risk early in life that resolves itself as they age. As they accumulate wealth through the lifecycle their precautionary savings motive falls, and they begin to save for approaching retirement.

FIGURE 4.10: DYNAMIC EFFECTS OF AN INCREASE IN MEANS-TESTED TRANSFERS TO WORKERS

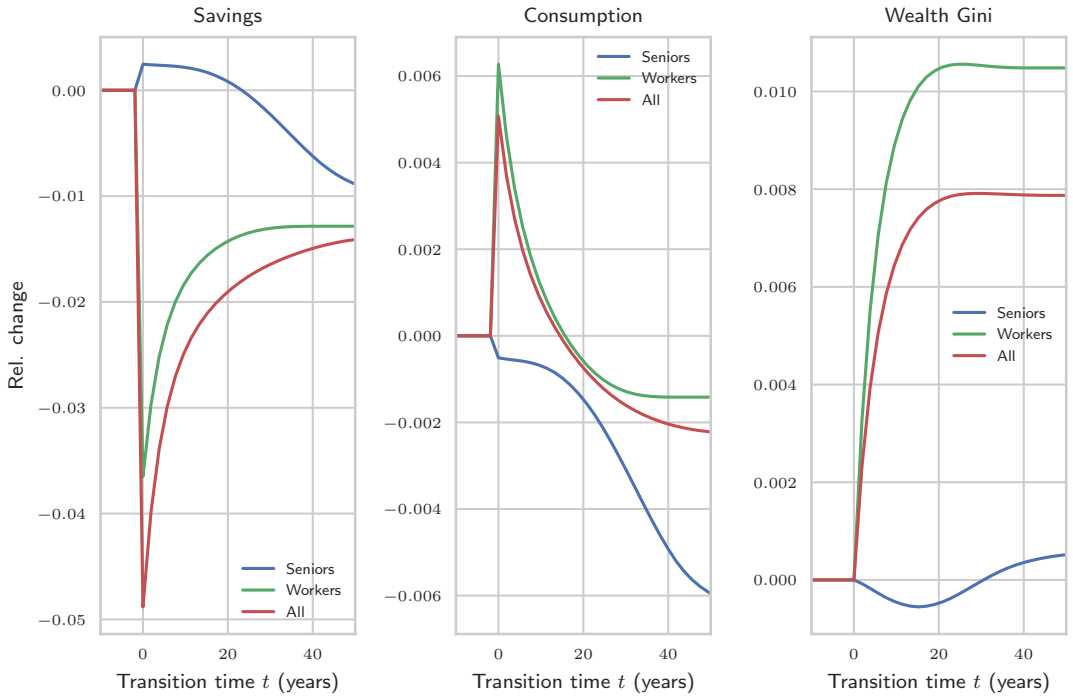


Figure 4.9 plots the change in long-run wealth inequality by age caused by various policies. Progressive transfers to workers increase wealth inequality among working-age households, by crowding out especially the savings of less-productive households. As households age, the precautionary savings motive weakens and this effect wanes. As a result, shocks to capital returns become relatively larger and churn the distribution of wealth, returning the inequality among seniors to levels comparable with the baseline.

Dynamics

Figure 4.10 plots the transitional effect of introducing the program of income-tested transfers to workers. In the short-run, workers find themselves able to increase consumption and reduce savings. Since many workers in the first generation are receiving a net windfall that will be financed by later generations, this effect is larger than the long-run change in consumption and savings.

The effect on seniors is weak at first, as they do not receive the transfers nor pay any tax. Eventually, however, the first generation of workers who have decreased their precautionary savings will reach age 65 and above. They will have accumulated less wealth on average during their working life, and so senior consumption begins to fall.

As we saw in Figure 4.10, the distributional effect of this reform is largest among workers. Indeed, the final panel shows that wealth inequality among workers shoots up over 20 years, driving an increase in overall wealth inequality, while the distribution of wealth among seniors is relatively unchanged.

The dynamics of increasing income-tested transfers to workers thus mirror those resulting from an increase in social security.

4.4 conclusion

In this paper, I build a framework for studying the distributional impact of fiscal policy changes. Using an overlapping generations heterogeneous agent model, I analyze the reactions of savings to changes in transfer programs, and quantify the resulting dynamics of the distribution of wealth. I find that an increase in social insurance leads to a long-run increase in wealth inequality, by crowding out the

savings of poorer agents. The distributional dynamics are slow, suggesting that expansions in U.S. social insurance programs in the 1960s may be associated with the rise in wealth inequality observed since the 1980s. Making transfers more progressive amplifies these effects. Subsidies to workers act in a similar way, crowding out the precautionary savings of lower-income households.

Further, these distributional dynamics highlight a potential vicious cycle for policymakers that have an explicit goal to reduce wealth inequality. Expanding transfer programs may reduce wealth inequality in the short-run, but increases it in the long-run, deceiving short-sighted policymakers.

To further develop this point, in forthcoming work I apply the framework of this paper to identify the determinants of this rise in U.S. wealth inequality. Economists debate whether this should be ascribed to fiscal factors or to structurally increasing labor income inequality. By modeling the evolution of U.S. postwar fiscal policy and income inequality, this approach can measure the relative contribution of each factor to wealth inequality. In light of this paper, it is likely that fiscal policy has played a larger role than previously thought, through the social insurance channel.

Another useful extension of this paper would be to examine changes to health risk and the extent to which health insurance affects household savings. Empirical evidence demonstrates that savings elasticities in this case are similar to those for social insurance. However, correlations between health outcomes and wealth and income would suggest that the distributional effects are large.

Table 4.2: Reactions of aggregates to pension reform size (elasticities)

Reform	Cons.	Saving	Wealth	Wealth Gini
Social security				
+10%	-0.01	-0.31	-0.20	0.03
-10%	-0.01	-0.33	-0.21	0.03
+1%	0.00	-0.31	-0.21	0.03
+100%	-0.01	-0.26	-0.17	0.02
Worker subsidies				
Small income test	-0.02	-0.13	-0.17	0.07
Larger income test	-0.04	-0.21	-0.29	0.13
Flat to sub-35s	0.00	-0.03	-0.01	0.00

4.A further analysis and robustness results

4.A.1 *Robustness of pension reform analysis*

In this section, I run further pension reform experiments, modifying the size, sign and shape of the reform relative to Section 4.3. Each of these is accompanied by an increase in the level α of the income tax rate in order to keep government budget balance.

First, in the first four rows of Table 4.2 I vary the size of the flat pension reform, reporting the elasticity of different variables to change in pension wealth. The reaction of aggregates and of inequality is quite linear in the size and sign of the change.

Next, I analyze different subsidies to workers. Each is calculated to cost the same to the government as a 10% increase in social security. Flat lump-sum transfers to households aged below 35 have negligible effect, as agents deduct the benefits from perceived future tax liability. The two income-tested transfer

schedules take the form

$$\Theta(t, w, z) = \max\{0, A(B - z)\}1_{t \leq 65}. \tag{4.6}$$

With this continuous formulation (in z), workers with the lowest productivity type receive the largest benefit. These taper off at rate A as productivity approaches the eligibility threshold B. The median productivity level is 1.0. The large income test has $(A, B) = (700, 0.5)$, while the smaller income test has $(A, B) = (1500, 1.0)$.

The two income-tested policies have a relatively smaller impact on aggregates than pension reforms. However, they have a large effect on wealth inequality; especially so for the more progressive policy. These are powerful insurative policies that crowd out the precautionary savings motive, argued by Gourinchas and Parker (2002) to account for the majority of savings of households aged below 40.

4.A.2 *Computation of human wealth and pension wealth*

The appropriate concept of human wealth in this model is given by,

$$h(t, w, z) = \mathbb{E}_t \left[\int_t^\infty (\Theta(t', w_{t'}, z_{t'}) + y_{t'} z_{t'}) e^{-\int_{t'}^{\xi} (r_s^b + \lambda_s)\, \mathrm{d}s}\, \mathrm{d}t' \mid w_t = w, z_t = z \right]. \tag{4.7}$$

Given the dynamics of z_t, we can compute h_t as follows. Define

$$H(t, w, z; \bar{t}) = \mathbb{E}_t \left[\int_t^{\bar{t}} (\Theta(t', w_{t'}, z_{t'}) + y_{t'} z_{t'}) e^{-\int_{t'}^{\xi} (r_s^b + \lambda_s)\, \mathrm{d}s}\, \mathrm{d}t' \mid w_t = w, z_t = z \right],$$

the expected discounted income from age t until some horizon $\bar{t}$ (so that h_t is the limit of H as $\bar{t} \to \infty$). Then H satisfies the Feynman–Kac equation[10]

$$\begin{aligned} 0 &= \partial_t H(t, w, z; \bar{t}) + \mathcal{A}H(t, w, z; \bar{t}) + (\Theta(t, w, z) + y_t z) - (r_t + \lambda_t)H(t, w, z; \bar{t}), \\ 0 &= H(\bar{t}, w, z; \bar{t}), \end{aligned} \tag{4.8}$$

where $\mathcal{A}$ is the infinitesimal generator for the joint process (w_t, z_t).

We can solve this equation numerically, working backwards in time from the terminal condition $0 = H(\bar{t}, w, z; \bar{t})$. In practice, this is done by splitting H into a component of human wealth at infinity (i.e. after the maximal age in the grid, after which the problem is stationary), and a lifecycle component. Solving the stationary analog of equation (4.8) yields the component at infinity. Paste this function as the terminal boundary condition at the maximal age, and solve backwards to get the lifecycle component.

Pension wealth

Denote pension wealth PW as $\tilde{h}$ here. It can be computed in the same way, by varying the payoff in equation (4.7) to include only social security transfers,

$$\tilde{h}(t, w, z) = \mathbb{E}_t \left[\int_t^\infty (\Theta(t', w_{t'}, z_{t'}) 1_{t' \geq 65}) e^{-\int_{t'}^{\xi} (r_s^b + \lambda_s)\, \mathrm{d}s}\, \mathrm{d}t' \mid w_t = w, z_t = z \right]. \tag{4.9}$$

Defining $\tilde{H}$ analogously, the appropriate Feynman–Kac condition becomes

$$0 = \partial_t \tilde{H}(t, w, z; \bar{t}) + \mathcal{A}\tilde{H}(t, w, z; \bar{t}) + (\Theta(t, w, z) 1_{t \geq 65}) - (r_t + \lambda_t)\tilde{H}(t, w, z; \bar{t}),$$

and the computational algorithm goes through as before.

[10] See, for example, Øksendal (2003).

CONCLUSION

In this thesis, I develop a framework for addressing questions in macroeconomics lying in the intersection of lifecycle and distributional considerations. I build tools for studying overlapping generations heterogeneous agent economies using new methods based on continuous time. This systematic approach, harnessing the similarities between the lifecycle and heterogeneity in income and wealth, strikes a balance between tractability and flexibility.

The first chapter sets out the framework and solution method in detail, with the goal of operationalizing the study of these economies in different contexts. I show that a simple version of the framework, relying solely on canonical mechanisms of lifecycle and precautionary savings behavior, yields empirically realistic distributions across age, income, and wealth.

The second chapter highlights the tractability of this framework, developing theoretical results on the shape and dynamics of the wealth distribution. I demonstrate that idiosyncratic risk to capital returns – a theory proposed to explain the fat tail in wealth observed in the U.S. – struggles to explain the changes in wealth inequality seen over the past decades. I show that this commonly used mechanism would explain current levels of wealth inequality by the presence of multicentenarian dynasties in the wealth tail, which were however not present in the 1970s.

The third chapter examines the distributional mechanisms of the model in

more detail, by examining three independent investigations. I quantify the contribution of lifecycle savings to total household wealth. Household savings vary in sign across the lifecycle, but I show they aggregate to a positive in a calibrated economy. As a result, savings due to lifecycle variation in income accounts for a 1.5 percentage point decrease in the risk-free rate relative to a model without lifecycle motives. Extending findings by Aiyagari (1994) and Huggett (1993) that idiosyncratic income risk can lead to oversaving, I thus show that lifecycle income variation further depresses safe rates. These results help explain and quantify the low prevailing safe interest rates, especially in countries with aging populations. These findings suggest that social security systems that smooth out lifecycle fluctuations in income could help boost the natural rate of interest in an aging population.

The second section of Chapter 3 analyzes two sources of secular stagnation in the model, population aging and rising income inequality. I show that both of these raise aggregate savings through the compositional channel: placing more wealth into the hands of households with higher savings rates. As a result, the risk-free rate falls with relatively high elasticity in the model. The final section of Chapter 3 examines the distributional impact of a permanent change to the risk-free interest rate. Common wisdom would say that borrowers would lose and lenders would gain from an increase in the risk-free rate. However, households are typically borrowers and lenders at different parts of their life cycle. As a result, an increase in borrowing rates has a scarring effect on almost all agents, so that even lenders tend to end up worse off in the long run. This exercise highlights the importance of considering both lifecycle and financial heterogeneity together.

The final chapter of this thesis quantifies the distributional impact of changes

to fiscal policy. Applying the framework of Chapter 1, I focus on the reaction of household savings to changes in social insurance, and calculate the resulting dynamic effects on the distribution of wealth. I demonstrate that an increase in social insurance can reduce short-run wealth inequality while *increasing* it in the long run, highlighting a vicious cycle for short-sighted policymakers. This is due to the crowding out of household savings by public transfers being strongest in the bottom half of the wealth distribution. The fact that these policies increase long-run wealth inequality despite being welfare improving highlights some of the difficulty in using wealth inequality as a proxy for social welfare.

Agenda for future research

Future research will extend and synthesize the findings of the chapters in this thesis. Firstly, the supply side of the models used in this thesis is stylized by design, focusing on household decisions and wealth. In future work, I plan to enrich the production sector in order to harness the channels set out in this thesis. Further, the framework of Chapter 1 is limited to idiosyncratic uncertainty. Recent results in the theory and methods of heterogeneous agent models[11] may help extend this framework to handle also aggregate risk. This would open the door to studying larger-scale dynamic stochastic general equilibrium models that incorporate lifecycle and financial heterogeneity. Forthcoming work with Guido Ascari and Andrea Ferrero uses such a model to study the effect of raising the inflation target in an aging population, embedding the households from Chapter 1 into a New Keynesian economy. Lifecycle and cross-sectional patterns in consumption, saving, and portfolio choice are crucial for pinning down the response

[11] These include Algan et al. (2014), Winberry (2016), and Ahn et al. (2018).

of aggregate variables and distributions to such a policy change. This paper will build upon the distributional results of Section 3.4.

Secondly, Chapter 4 indicates that social insurance can lead to more wealth inequality, with a lag of up to twenty years. There is an ongoing debate on the determinants of the rise in wealth inequality in the United States. Previous studies have argued about whether this should be ascribed to declining marginal tax rates or to increasing labor income inequality due to skill-biased technological change. The results of Chapter 4 suggest that there may be a connection between the dramatic growth in U.S. government transfer programs from the 1960s and rising wealth inequality in the U.S. from the 1980s. In a companion paper, I address this issue by modeling the changes in the tax and transfer system in the postwar United States, as well as changes in the distribution of income. By performing counterfactual experiments in the vein of Kaymak and Poschke (2016), I plan to quantify the relative contribution to wealth inequality of taxes, transfers, and structural changes to income over that period. Previous studies have weighed tax against structural changes, finding a larger role for changes to income. However, Chapter 4 and preliminary results suggest that in such a case the role of fiscal policy may be understated, due to a large reaction of wealth inequality to changes in social insurance. This analysis would help address questions raised by Gabaix et al. (2016) and Chapter 2, which highlight the difficulty of common mechanisms to explain the dynamics of U.S. wealth inequality over the past decades.

BIBLIOGRAPHY

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L Violante. 2013. *Education policy and intergenerational transfers in equilibrium.* Technical report. National Bureau of Economic Research.

Achdou, Yves, Francisco J Buera, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2014. "Partial differential equation models in macroeconomics." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 372 (2028): 20130397.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2017. *Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach.* Technical report. National Bureau of Economic Research.

Aguila, Emma. 2011. "Personal retirement accounts and saving." *American Economic Journal: Economic Policy* 3 (4): 1–24.

Ahn, SeHyoun, Greg Kaplan, Benjamin Moll, Thomas Winberry, and Christian Wolf. 2018. "When inequality matters for macro and macro matters for inequality." *NBER Macroeconomics Annual* 32 (1): 1–75.

Aiyagari, S Rao. 1994. "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics:* 659–684.

Algan, Yann, Olivier Allais, Wouter J Den Haan, and Pontus Rendahl. 2014. "Solving and simulating models with heterogeneous agents and aggregate uncertainty." In *Handbook of Computational Economics,* 3:277–324. Elsevier.

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman. 2017. *Tax evasion and inequality.* Technical report. National Bureau of Economic Research.

Ameriks, John, and Stephen P Zeldes. 2004. *How do household portfolio shares vary with age.* Technical report.

Ampudia, Miguel, Dimitris Georgarakos, Jiri Slacalek, Oreste Tristani, Philip Vermeulen, and Giovanni Violante. 2018. "Monetary Policy and Household Inequality."

Ascari, Guido, Andrea Ferrero, and Pierre E. Weil. 2018. "Distributional effects of a change in the inflation target."

Attanasio, Orazio P, and Agar Brugiavini. 2003. "Social security and households' saving." *The Quarterly Journal of Economics* 118 (3): 1075–1119.

Attanasio, Orazio P, and Susann Rohwedder. 2003. "Pension wealth and household saving: Evidence from pension reforms in the United Kingdom." *American Economic Review* 93 (5): 1499–1521.

Auclert, Adrien. 2017. *Monetary policy and the redistribution channel.* Technical report. National Bureau of Economic Research.

Auerbach, Alan J, and Laurence J Kotlikoff. 1985. "Simulating Alternative Social Security Responses to the Demographic Transition." *National Tax Journal* 38 (2): 153.

Auerbach, Alan J, Laurence J Kotlikoff, et al. 1987. *Dynamic fiscal policy.* Cambridge University Press.

Bardi, M., and I. Capuzzo-Dolcetta. 2008. *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman Equations.* Modern Birkhäuser Classics. Birkhäuser Boston. isbn: 9780817647544.

Barles, Guy, and Panagiotis E Souganidis. 1991. "Convergence of approximation schemes for fully nonlinear second order equations." *Asymptotic Analysis* 4 (3): 271–283.

Bénabou, Roland. 2000. "Unequal societies: Income distribution and the social contract." *American Economic Review* 90 (1): 96–129.

Benhabib, Jess, Alberto Bisin, and Mi Luo. 2019. "Wealth distribution and social mobility in the US: A quantitative approach." *American Economic Review* 109 (5): 1623–47.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu. 2011. "The distribution of wealth and fiscal policy in economies with finitely lived agents." *Econometrica* 79 (1): 123–157.

———. 2015. "The wealth distribution in Bewley economies with capital income risk." *Journal of Economic Theory* 159:489–515.

———. 2016. "The distribution of wealth in the Blanchard–Yaari model." *Macroeconomic Dynamics* 20 (02): 466–481.

Bertrand, Marianne, and Adair Morse. 2016. "Trickle-down consumption." *Review of Economics and Statistics* 98 (5): 863–879.

Bewley, Truman. 1986. "Stationary monetary equilibrium with a continuum of independently fluctuating consumers." *Contributions to mathematical economics in honor of Gérard Debreu* 79.

Blanchard, Olivier J. 1985. "Debt, deficits, and finite horizons." *Journal of Political Economy* 93 (2): 223–247.

Blundell, Richard, Martin Browning, and Costas Meghir. 1994. "Consumer demand and the life-cycle allocation of household expenditures." *The Review of Economic Studies* 61 (1): 57–80.

Bofinger, Peter, and Philipp Scheuermeyer. 2016. "Income Distribution and Aggregate Saving: A Non-Monotonic Relationship." *Review of Income and Wealth.*

Bottazzi, Renata, Tullio Jappelli, and Mario Padula. 2006. "Retirement expectations, pension reforms, and their impact on private wealth accumulation." *Journal of Public Economics* 90 (12): 2187–2212.

Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus. 2016. "Estimating Top Income and Wealth Shares: Sensitivity to Data and Methods." *American Economic Review* 106, no. 5 (May): 641–45. doi:10.1257/aer.p20161020. http://www.aeaweb.org/articles?id=10.1257/aer.p20161020.

Brown, Jeffrey R. 2007. *Rational and behavioral perspectives on the role of annuities in retirement planning.* Technical report. National Bureau of Economic Research.

Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico. 2018. "The consumption response to positive and negative income shocks." *Journal of Monetary Economics* 96:1–15.

Bütler, Monika. 2001. "Neoclassical life-cycle consumption: A textbook example." *Economic Theory* 17 (1): 209–221.

Candler, Graham V. 1998. *Finite-Difference Methods for Continuous-Time Dynamic Programming.* QM&RBC Codes, Quantitative Macroeconomics & Real Business Cycles. https://ideas.repec.org/c/dge/qmrbcd/129.html.

Capuzzo–Dolcetta, Italo, and Pierre–Louis Lions. 1990. "Hamilton–Jacobi equations with state constraints." *Transactions of the American Mathematical Society* 318 (2): 643–683.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N White. 2017. "The distribution of wealth and the marginal propensity to consume." *Quantitative Economics* 8 (3): 977–1020.

Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio. 2016. "Demographics and real interest rates: inspecting the mechanism." *European Economic Review.*

Castaneda, Ana, Javier Diaz-Gimenez, and Jose-Victor Rios-Rull. 2003. "Accounting for the US earnings and wealth inequality." *Journal of Political Economy* 111 (4): 818–857.

Champernowne, David G. 1953. "A model of income distribution." *The Economic Journal* 63 (250): 318–351.

Chetty, Raj. 2015. "Behavioral economics and public policy: A pragmatic perspective." *American Economic Review* 105 (5): 1–33.

Chetty, Raj, John N Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen. 2014. "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark." *The Quarterly Journal of Economics* 129 (3): 1141–1219.

Cocco, Joao F, Francisco J Gomes, and Pascal J Maenhout. 2005. "Consumption and portfolio choice over the life cycle." *The Review of Financial Studies* 18 (2): 491–533.

Conesa, Juan C, and Dirk Krueger. 1999. "Social security reform with heterogeneous agents." *Review of Economic Dynamics* 2 (4): 757–795.

De Nardi, Mariacristina, Eric French, and John Bailey Jones. 2006. *Differential mortality, uncertain medical expenses, and the saving of elderly singles.* Technical report. National Bureau of Economic Research.

———. 2016. "Medicaid insurance in old age." *American Economic Review* 106 (11): 3480–3520.

Deaton, Angus. 2016. "Measuring and understanding behavior, welfare, and poverty." *American Economic Review* 106 (6): 1221–43.

Deaton, Angus S, Pierre-Olivier Gourinchas, and Christina Paxson. 2002. "Social security and inequality over the life cycle." In *The distributional aspects of Social Security and Social Security reform,* 115–148. University of Chicago Press.

Deaton, Angus, and Christina Paxson. 1994. "Intertemporal choice and inequality." *Journal of Political Economy* 102 (3): 437–467.

Dixit, Avinash K. 1993. *The Art of Smooth Pasting.* Vol. 55. Taylor & Francis.

Dynan, Karen E, Jonathan Skinner, and Stephen P Zeldes. 2004. "Do the rich save more?" *Journal of Political Economy* 112 (2): 397–444.

Economic Co-Operation, Organisation for, and Development. 2015. *Pensions at a Glance 2015: OECD and G20 Indicators.* OECD pensions at a glance. OECD. isbn: 9789264240636. https://books.google.co.uk/books?id=p3fojgEACAAJ.

Eggertsson, Gauti B, and Neil R Mehrotra. 2014. *A model of secular stagnation.* Technical report. National Bureau of Economic Research.

Fagereng, Andreas, Martin Blomhoff Holm, and Gisle James Natvik. 2016. *MPC heterogeneity and household balance sheets.* Technical report.

Feigenbaum, James. 2008. "Can mortality risk explain the consumption hump?" *Journal of Macroeconomics* 30 (3): 844–872.

Feng, Jin, Lixin He, and Hiroshi Sato. 2011. "Public pension and household saving: Evidence from urban China." *Journal of Comparative Economics* 39 (4): 470–485.

Fernández-Villaverde, Jesús, and Dirk Krueger. 2007. "Consumption over the life cycle: Facts from consumer expenditure survey data." *The Review of Economics and Statistics* 89 (3): 552–565.

Fisher, Irving. 1930. *Theory of interest: as determined by impatience to spend income and opportunity to invest it.* Augustus Kelly Publishers, Clifton.

Friedman, Benjamin M, and Mark J Warshawsky. 1990. "The cost of annuities: Implications for saving behavior and bequests." *The Quarterly Journal of Economics* 105 (1): 135–154.

Gabaix, Xavier. 2009. "Power Laws in Economics and Finance." *Annual Review of Economics* 1 (1): 255–294. doi:10.1146/annurev.economics.050708.142940.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. "The dynamics of inequality." *Econometrica* 84 (6): 2071–2111.

Gertler, Mark. 1999. "Government debt and social security in a life-cycle economy." In *Carnegie-Rochester Conference Series on Public Policy,* 50:61–110. Elsevier.

Gourinchas, Pierre-Olivier, and Jonathan A Parker. 2002. "Consumption over the life cycle." *Econometrica* 70 (1): 47–89.

Gruber, Jonathan, and Aaron Yelowitz. 1999. "Public health insurance and private savings." *Journal of Political Economy* 107 (6): 1249–1274.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2015. *What do data on millions of US workers reveal about life-cycle earnings risk?* Technical report. National Bureau of Economic Research.

Hansen, Alvin H. 1939. "Economic progress and declining population growth." *The American Economic Review* 29 (1): 1–15.

Hansen, Gary D, and Selahattin İmrohoroğlu. 2008. "Consumption over the life cycle: The role of annuities." *Review of Economic Dynamics* 11 (3): 566–583.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2009. "Quantitative macroeconomics with heterogeneous households." *Annu. Rev. Econ.* 1 (1): 319–354.

———. 2010. "The macroeconomic implications of rising wage inequality in the United States." *Journal of political economy* 118 (4): 681–722.

———. 2014. "Consumption and labor supply with partial insurance: An analytical framework." *American Economic Review* 104 (7): 2075–2126.

———. 2017. "Optimal tax progressivity: An analytical framework." *The Quarterly Journal of Economics* 132 (4): 1693–1754.

———. 2019. *Optimal progressivity with age-dependent taxation.* Technical report. National Bureau of Economic Research.

Hubbard, R Glenn, Jonathan Skinner, and Stephen P Zeldes. 1995. "Precautionary saving and social insurance." *Journal of Political Economy* 103 (2): 360–399.

Hubmer, Joachim, Per Krusell, and Anthony A Smith Jr. 2016. *The historical evolution of the wealth distribution: A quantitative-theoretic investigation.* Technical report. National Bureau of Economic Research.

Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of Economic Dynamics and Control* 17 (5): 953–969.

———. 1996. "Wealth distribution in life-cycle economies." *Journal of Monetary Economics* 38 (3): 469–494.

Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of lifetime inequality." *American Economic Review* 101 (7): 2923–54.

Hurst, Erik. 2008. *The retirement of a consumption puzzle.* Technical report. National Bureau of Economic Research.

İmrohoroglu, Ayşe, Selahattin Imrohoroglu, and Douglas H Joines. 1995. "A life cycle analysis of social security." *Economic theory* 6 (1): 83–114.

International Monetary Fund. 2017. *Fiscal Monitor.*

Japanese National Institute of Population and Social Security Research. 2017. *Japanese Mortality Database.* Accessed: 2017-01-01. `http://www.ipss.go.jp/p-toukei/JMD/index-en.asp`.

Kaplan, Greg. 2012. "Inequality and the life cycle." *Quantitative Economics* 3 (3): 471–525.

Kaplan, Greg, Benjamin Moll, and Giovanni L Violante. 2018. "Monetary policy according to HANK." *American Economic Review* 108 (3): 697–743.

Kaplan, Greg, and Giovanni L Violante. 2010. "How much consumption insurance beyond self-insurance?" *American Economic Journal: Macroeconomics* 2 (4): 53–87.

———. 2014. "A model of the consumption response to fiscal stimulus payments." *Econometrica* 82 (4): 1199–1239.

———. 2018. "Microeconomic heterogeneity and macroeconomic shocks." *Journal of Economic Perspectives* 32 (3): 167–94.

Kaplan, Greg, Giovanni L Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity* 2014 (1): 77–138.

Karatzas, Ioannis, and Steven E. Shreve. 1991. *Brownian Motion and Stochastic Calculus.* Graduate Texts in Mathematics. Springer New York. isbn: 9780387976556.

Kartashova, Katya. 2014. "Private Equity Premium Puzzle Revisited." *American Economic Review* 104 (10): 3297–3334.

Kaymak, Barış, and Markus Poschke. 2016. "The evolution of wealth inequality over half a century: The role of taxes, transfers and technology." *Journal of Monetary Economics* 77:1–25.

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money.* New York: Harcourt, Brace.

Kindermann, Fabian, and Dirk Krueger. 2015. *High marginal tax rates on the top 1%? Lessons from a life cycle model with idiosyncratic income risk.* Technical report. National Bureau of Economic Research.

Kopczuk, Wojciech. 2015. "What do we know about the evolution of top wealth shares in the United States?" *The Journal of Economic Perspectives* 29 (1): 47–66.

Krueger, Dirk, and Alexander 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.

Kuhn, Moritz, and José-Victor Rios-Rull. 2016. "2013 Update on the US earnings, income, and wealth distributional facts: A View from Macroeconomics." *Quarterly Review, Federal Reserve Bank of Minneapolis, April:* 1–75.

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins. 2017. *Income and Wealth Inequality in America, 1949-2013.* Technical report. CESifo Working Paper.

Lachowska, Marta, and Michał Myck. 2018. "The effect of public pension wealth on saving and expenditure." *American Economic Journal: Economic Policy* 10 (3): 284–308.

Lasry, Jean-Michel, and Pierre-Louis Lions. 2006. "Jeux à champ moyen. I–Le cas stationnaire." *Comptes Rendus Mathématique* 343 (9): 619–625.

———. 2007. "Mean field games." *Japanese Journal of Mathematics* 2 (1): 229–260.

Lindert, Peter H. 2017. *The Rise and Future of Progressive Redistribution.* Technical report. Tulane University, Department of Economics.

Luttmer, Erzo GJ. 2011. "On the mechanics of firm growth." *The Review of Economic Studies* 78 (3): 1042–1068.

Maynard, Alex, and Jiaping Qiu. 2009. "Public insurance and private savings: who is affected and by how much?" *Journal of Applied Econometrics* 24 (2): 282–308.

Merton, Robert C. 1969. "Lifetime portfolio selection under uncertainty: The continuous-time case." *The Review of Economics and Statistics:* 247–257.

Misra, Kanishka, and Paolo Surico. 2014. "Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs." *American Economic Journal: Macroeconomics* 6 (4): 84–106.

Mitchell, Olivia S, James M Poterba, Mark J Warshawsky, and Jeffrey R Brown. 1999. "New evidence on the money's worth of individual annuities." *American Economic Review* 89 (5): 1299–1318.

Moskowitz, Tobias J, and Annette Vissing-Jørgensen. 2002. "The returns to entrepreneurial investment: A private equity premium puzzle?" *American Economic Review* 92 (4): 745–778.

Øksendal, Bernt. 2003. *Stochastic Differential Equations: An Introduction with Applications.* Springer Science & Business Media.

Olafsson, Arna, and Michaela Pagel. 2018. *The Retirement-Consumption Puzzle: New Evidence from Personal Finances.* Technical report. National Bureau of Economic Research.

Pareto, Vilfredo. 1896. *Cours d'économie politique.* Vol. 1. Librairie Droz.

Pham-Dao, Lien. 2016. *Public insurance and wealth inequality-a euro area analysis.* Technical report.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *The Quarterly Journal of Economics* 118 (1): 1–41.

———. 2007. "How progressive is the US federal tax system? A historical and international perspective." *Journal of Economic Perspectives* 21 (1): 3–24.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2017. "Distributional national accounts: methods and estimates for the United States." *The Quarterly Journal of Economics* 133 (2): 553–609.

Ríos-Rull, José-Víctor. 1995. "Models with heterogeneous agents." *Frontiers of business cycle research:* 98–125.

Saez, Emmanuel, and Gabriel 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.

Saporta, Benoite de. 2004. "Etude de la solution stationnaire de l'équation $Y_{n+1} = a_n Y_n + b_n$ à coefficients aléatoires." PhD diss., Rennes 1.

———. 2005. "Tail of the stationary solution of the stochastic equation $Y_{n+1} = a_n Y_n + b_n$ with Markovian coefficients." *Stochastic Processes and their Applications* 115 (12): 1954–1978.

Soner, Halil Mete. 1986. "Optimal control with state-space constraint II." *SIAM Journal on Control and Optimization* 24 (6): 1110–1122.

Sterk, Vincent, and Silvana Tenreyro. 2018. "The transmission of monetary policy through redistributions and durable purchases." *Journal of Monetary Economics* 99:124–137.

Storesletten, Kjetil, Christopher I Telmer, and Amir Yaron. 2004. "Consumption and risk sharing over the life cycle." *Journal of Monetary Economics* 51 (3): 609–633.

Straub, Ludwig. 2017. "Consumption, Savings, and the Distribution of Permanent Income." *Manuscript, November.*

Summers, Lawrence H. 2015. "Demand side secular stagnation." *American Economic Review* 105 (5): 60–65.

Thaler, Richard H, and Shlomo Benartzi. 2004. "Save more tomorrow™: Using behavioral economics to increase employee saving." *Journal of Political Economy* 112 (S1): S164–S187.

Thurow, Lester C. 1969. "The optimum lifetime distribution of consumption expenditures." *The American Economic Review* 59 (3): 324–330.

United Nations. 2015. *World Population Prospects: The 2015 Revision.*

Weil, Philippe. 1989. "The equity premium puzzle and the risk-free rate puzzle." *Journal of Monetary Economics* 24 (3): 401–421.

Winberry, Thomas. 2016. *A toolbox for solving and estimating heterogeneous agent macro models.* Technical report.

Wong, Arlene. 2015. *Population Aging and the Transmission of Monetary Policy to Consumption.* Technical report.

Yaari, Menahem E. 1965. "Uncertain lifetime, life insurance, and the theory of the consumer." *The Review of Economic Studies* 32 (2): 137–150.

Ziliak, James P. 2003. "Income transfers and assets of the poor." *Review of Economics and Statistics* 85 (1): 63–76.