TESI DOCTORAL UPF / 2021

Inês Martins Xavier

Essays in Macro-Finance and Wealth Inequality

TESI DOCTORAL UPF / 2021

Universitat
Pompeu Fabra
Barcelona

Essays in Macro-Finance and Wealth Inequality

Inês Martins Xavier

Essays in Macro-Finance and Wealth Inequality

Inês Xavier

TESI DOCTORAL UPF / 2021

Directors de la tesi

Fernando Broner and Alberto Martin

Department d'Economia i Empresa

Universitat Pompeu Fabra Barcelona

Aos meus pais

Acknowledgements

I want to express my deep gratitude to my advisors Fernando Broner and Alberto Martin. This thesis would not have been possible without their encouragement and guidance. By combining profound knowledge, kind patience and great sense of humour, they have set an example I hope to match one day.

I am grateful to Jan Eeckhout who helped me a lot during the Job Market, not only with his letter, but also with extensive comments on my paper and with great encouragement. Luca Fornaro, Jordi Galí, Dmitry Kuvshinov, Benjamin Moll, Claudia Sahm, Edouard Schaal and Jaume Ventura helped me with insightful feedback on different parts of this thesis. More broadly, I have greatly benefited from the extensive macro group at CREI and UPF with whom I shared many seminars and learned from. A special thanks to Marta Araque for the reliable and incredibly efficient support to all PhD students, including me.

Among the many wonderful people that I met at UPF over the years, I'd like to single out Juan, Eva, Francisco, Stefanie, Yiru, Chris, Ilja, Adil, Thomas, Milena, Kinga, Rasmus and Elisa, who made the PhD journey less lonely and more fun. Rita, Filipa, Mariana and Luca supported me every step of the way and always welcomed me back home with open arms. My Nunca Unca friends became more than my football team and helped me feel at home in Barcelona. Soto and Karo supported me during a virtual Job Market in a small apartment and in the midst of a pandemic. I am grateful for their patience. Teresa, Júlia, Matt, Joan, Tia Anna, Quim and Miquel taught me Catalan and generosity. I thank them for welcoming me to the family without reservations. The kind encouragement of Luís Delfim was determinant for my decision to join Universitat Pompeu Fabra, for which I am grateful.

The support of my family has been instrumental, not only for this thesis, but for navigating life in general. I am grateful for my mom whose strength and perseverance has always been one of my greatest inspirations. My dad, who swapped tennis courts for airports, has always been by my side. I thank my sister Joana for making the world a little easier for her younger siblings. My twin sister Marta – the only doctor whom anyone will call in case of a true emergency – has probably shared more time with me than anyone else in the world. Fati and Zé, the first economists that I ever met, showed me their support in many ways, including some delicious pães de queijo e fiambre that tasted like home.

Last, but not least, I am immensely grateful to Anna who joined me on this journey during the early stages of this thesis. I am thankful for her unconditional support and for making it all the more fun.

Abstract

This thesis investigates how asset markets affect macroeconomic outcomes and inequality. In the first chapter, I study the contribution of return heterogeneity for wealth inequality. First, I document that wealthier U.S. households earn, on average, higher returns on their wealth: moving from the 20th to the 99th percentile of the wealth distribution raises the average return on wealth from 3.6% to 8.3% a year. Second, I find that considering both earnings and return heterogeneity in a model of household wealth accumulation can fully account for the top 10% wealth share observed in the data (76%), which cannot be explained by earnings differences alone. The second chapter is dedicated to the study of asset price bubbles in economies that are vulnerable to a secular stagnation. I show that, when liquidity is scarce, bubbles can be useful as they expand the supply of assets and absorb resources that would otherwise be wasted. In doing so, bubbles crowd in aggregate consumption and raise welfare. The greater is the risk of collapsing, the lower is the likelihood that bubbles avoid stagnation. However, some risk may still be desirable to the extent that it relaxes financing constraints.

Resumen

Esta tesis investiga cómo los mercados de activos afectan la actividad económica y la desigualdad. En el primer capítulo, estudio la contribución de la heterogeneidad de retorno para la desigualdad en la distribución de la riqueza. En primer lugar, documento que los hogares estadounidenses más ricos obtienen, en promedio, mayores retornos de su riqueza: pasar del percentil 20 al 99 de la distribución de la riqueza implica un aumento de la tasa de retorno anual del 3.6% al 8.3%. En segundo lugar, demuestro que considerar tanto la heterogeneidad de los ingresos laborales como la del retorno del capital permite explicar que el 10% de la población más rica de Estados Unidos posee el 76% de la riqueza total del país, lo que no puede explicarse únicamente por las diferencias de ingresos. El segundo capítulo está dedicado al estudio de las burbujas en los precios de los activos en economías vulnerables a un estancamiento secular. Demuestro que, cuando la liquidez es escasa, las burbujas pueden ser útiles ya que amplían la oferta de activos y absorben recursos que de otro modo se desperdiciarían. Al hacerlo, las burbujas aumentan el consumo agregado y el bienestar. Cuanto mayor es el riesgo de colapso, menor es la probabilidad de que las burbujas puedan evitar el estancamiento. Sin embargo, puede ser deseable cierto nivel de riesgo ya que relaja restricciones financieras.

Preface

This doctoral thesis brings together two self-contained essays that study the aggregate and the distributional implications of financial markets broadly defined.

In the first chapter, I investigate the contribution of return heterogeneity for wealth inequality. This work joins a growing macro-inequality literature that investigates the reasons behind the high concentration of wealth observed in several countries over different periods of time. I focus on the United States which provides a stark example of wealth concentration: in 2019, just 10% of the families owned 76% of the economy's total wealth. I make two contributions. First, I use household-level data from the U.S. Survey of Consumer Finances (1989–2019) to investigate the degree of heterogeneity in returns to wealth. I find that wealthier households earn, on average, higher returns: moving from the 20th to the 99th percentile of the wealth distribution raises the average yearly return from 3.6% to 8.3%. Second, I study the implications of return differences for the distribution of wealth by incorporating realistic return heterogeneity into the workhorse model of earnings inequality. This exercise suggests that considering both earnings and return heterogeneity can fully account for the top 10% wealth share observed in the data (76%), while earnings differences can only explain about half of it. Overall, this paper provides evidence that return heterogeneity is crucial to understand top wealth shares in the United States.

In the second chapter, I study the effects of asset price bubbles in low interest rate economies. Two observations motivate this project. On the one hand, real interest rates have declined steadily over the past three decades in much of the western world, contributing to the revival of the secular stagnation hypothesis: the idea that structural factors may create a chronic excess of savings relative to the demand for new investments, depressing interest rates, output and growth. On the other hand, the last thirty years have been marked by recurrent episodes of large fluctuations in asset prices, often associated with "bubbles" due to an apparent disconnection between asset prices and fundamentals. In this paper, I develop a model to understand the implications of bubbles in low interest rate environments. Three broad insights emerge from the analysis. First, bubbles have different implications in good and in bad times. Outside of the stagnation environment, bubbles crowd out private lending and reduce welfare by tightening borrowing constraints and distorting consumption decisions. On the contrary, bubbles can be expansionary and raise welfare in times of stagnation, when aggregate demand is chronically low and the economy operates below full capacity. In this case, bubbles absorb resources that would otherwise be wasted and crowd in aggregate consumption. Second, bubbles make stagnation less likely. On the one hand, the stagnation regime becomes less likely as bubbles raise the natural interest rate and can prevent output from falling below potential. On the other hand, bubbles provide a mechanism to get out of stagnation by expanding aggregate demand and raising equilibrium employment. Finally, bubbles that may randomly collapse are less likely to avoid stagnation than safe bubbles. However, some bubble risk may be desirable provided it is not too high. Riskier bubbles are associated with lower interest rates, which loosen borrowing constraints and may improve the allocation of resources.

Contents

List of Figures

List of Tables

Chapter 1

Wealth Inequality in the US: the Role of Heterogeneous Returns

1.1 Introduction

In the United States, wealth is highly concentrated. The richest 10% households own 76% of the economy's total wealth, half of which is actually owned by the wealthiest 1%.[1] Understanding what produces highly skewed wealth distributions is the goal of a growing macro-inequality literature making use of increasingly rich micro datasets.

The largest strand of the literature has focused primarily on labor income differences as the chief driver of wealth dispersion. However, this work has tended to conclude that, although realistic and important, differences in labor income are not enough to explain the large wealth concentration observed in the data (De Nardi and Fella, 2017)[2]. This follows from the inability of the workhorse model of labor income heterogeneity to reproduce the high top wealth shares observed in the data. In particular, this class of models predicts that wealth cannot be more concentrated than earnings which is at odds with the empirical evidence that top wealth shares are larger, and decay slower, than top earnings shares[3].

The failure of the workhorse model to explain high wealth concentration has prompted researchers to look for other potential explanations. Return heterogeneity is one of the proposed mechanisms[4]. Theoretically, it has been shown that return heterogeneity is not only a powerful force towards wealth concentration, but it can potentially explain the empirical fact that top wealth shares are considerably larger than top earnings shares (Benhabib and Bisin, 2018). However, the lack of empirical evidence on U.S. wealth returns has made it difficult to ascertain

[1]Based on the 2019 U.S. Survey of Consumer Finances.

[2]De Nardi and Fella (2017) provide a broad overview of this literature and its main results.

[3]In this class of models, the right tail of the wealth distribution cannot be thicker than that of the earnings distribution. Section 1.2 provides empirical evidence on the distributions of earnings and wealth in the United States.

[4]Benhabib et al. (2011, 2015), Nirei and Aoki (2016) and Gabaix et al. (2016) were some of the first papers to consider return heterogeneity in micro-founded models of consumption and savings.

their contribution for wealth inequality.

In this paper, I make two contributions. First, I use household-level data from the U.S. Survey of Consumer Finances (SCF) from 1989 to 2019 to investigate the degree of heterogeneity in wealth returns. Second, I study the implications of return heterogeneity for the distribution of wealth in the United States by incorporating return heterogeneity into the workhorse model of earnings risk, and calibrating returns to match the empirical evidence on returns to wealth.

The analysis of U.S. household data uncovers the following empirical patterns. First, the average return on wealth is increasing in households' wealth, in line with the findings of Bach et al. (2020) and Fagereng et al. (2020) for Sweden and Norway, respectively. Moving from the 20th to the 99th percentile of the wealth distribution raises the average yearly return on wealth by 4.7 percentage points, from 3.6% to 8.3%.[5] One source of the observed differentials is the allocation of wealth between the different asset classes. The asset portfolios at the top of the wealth distribution tend to own a larger share of equity than bottom or middle ones, in which residential real estate predominates. This generates heterogeneity in wealth returns because equity earns, on average, higher returns than real estate: a premium of 8.1 and 1.4 percentage points for private businesses and public equity, respectively. Moreover, this paper uncovers a second source of heterogeneity between households. Even within narrow asset classes, returns exhibit important differences and tend to increase with wealth. This is particularly true in the case of private businesses and real estate.

To understand the quantitative implications of the estimated return differences, I develop a partial equilibrium model of household saving behavior. Individuals face two sources of heterogeneity: earnings and the rate of return on wealth. Given the exogenous processes for earnings and returns, they optimally decide how much to save and consume at each point in time. I consider a standard process for the individual earnings process. Households face a stochastic income process and receive different realizations of a persistent earnings shock which creates dispersion in savings and wealth. Turning to return heterogeneity, I consider two mechanisms that drive return differences. The first one is analogous to earnings: ex-post *luck* that makes some individuals earn higher returns than others. The second mechanism – *type dependence* – allows for persistent ex-ante differences among individuals. Specifically, individuals may face fundamentally different return processes (e.g. different mean return), and the "high return types" end up accumulating more wealth.

The main result of the quantitative exercise is that return heterogeneity, calibrated to the U.S. economy, generates a considerable amount of wealth concentration. I show that a simple model that accounts for return heterogeneity, in addition to earnings inequality, can fully account for the top 10% wealth share observed in the data (76%). To get a sense of the importance of return heterogeneity, I estimate a counterfactual model economy in which I shut down all return differences, and earnings heterogeneity is the only source of wealth dispersion. The top

[5] I present evidence for the positive domain of the wealth distribution. As the bottom 20% of the U.S. wealth distribution have negative wealth, I omit this group to avoid confusion when interpreting the estimates.

10% wealth share implied by such model is equal to 36%, implying much less wealth concentration than in the data. This exercise suggests that return heterogeneity is at least as important as labor income differences to understand wealth inequality in the United States, in particular the distribution of wealth between the top 10% and the remaining 90% households. Within the top 10% of the wealth distribution, the calibrated model of return heterogeneity implies that the top 5% and the top 1% own 69% and 55% of the total wealth, respectively. These numbers are somewhat larger than their empirical counterparts (65% and 37%), implying a slightly thicker distribution tail. Importantly, they imply much more wealth concentration than the workhorse model of earnings heterogeneity (21% and 5%, respectively). Overall, these findings suggest that return heterogeneity is crucial to understand top wealth shares in the United States.

Related literature. This work relates to several strands of the literature on wealth inequality. First, it relates to the recent literature that uses disaggregated micro-data to estimate wealth returns. In two key contributions, Bach et al. (2020) and Fagereng et al. (2020) document a substantial degree of heterogeneity in individual wealth returns based on Swedish and Norwegian administrative data. Importantly, both papers find a positive correlation between wealth and returns. This paper contributes to this literature by showing evidence of return heterogeneity in the United States, which differs in several ways from Scandinavian economies, including in the degree of wealth inequality[6]. In the absence of administrative data on the asset holdings of U.S. households, I propose a methodology based on survey data. Despite the accessibility and popularity of the U.S. Survey of Consumer Finances, there is no previous research documenting heterogeneity in U.S. household returns and, in particular, how they vary with wealth. To calculate the return on public and private equity, I build on Moskowitz and Vissing-Jørgensen (2002) and Kartashova (2014) who use data from the SCF to calculate the return on these asset classes from 1990 to 2010. However, these papers focus on the economy-wide return and do not investigate whether there is heterogeneity between households. Additionally, I estimate the return on the remaining wealth components included in the survey (e.g., deposits, bonds and real estate). Finally, this paper closely relates to Kuhn et al. (2020) who highlight the importance of asset price fluctuations and household portfolio heterogeneity for wealth dynamics in postwar America.

On the theoretical side, this paper relates to the macro-inequality literature that uses models with different sources of heterogeneity to understand the dynamics and distribution of wealth. Bewley (1977), İmrohoroğlu (1992), Huggett (1993) and Aiyagari (1994) provided the first contributions to what has become the workhorse macroeconomic model to study wealth inequality based on labor income heterogeneity. As previously discussed, this class of models implies too little wealth concentration compared to the data, which has motivated several extensions that aim to match the data better. Closest to this paper is the strand of the literature

[6]The top 10% wealth share is close to 52% in Sweden (calculation based on Bach et al. (2020)). In Norway in 2009, the top 10% wealth share was equal to 53% (Epland and Kirkeberg, 2012).

highlighting the implications of capital return heterogeneity. Quadrini (2000) and Cagetti and De Nardi (2006) explicitly consider idiosyncratic returns to entrepreneurship and show that this mechanism can generate a thick wealth distribution tail. Relatedly, Benhabib et al. (2011) show analytically that idiosyncratic return risk can generate a Pareto tailed wealth distribution whose thickness is driven by the heterogeneity in returns to wealth and not to human capital. Gabaix et al. (2016) explore the dynamics of wealth over time and suggest that in order to generate fast changes in tail inequality in the magnitude measured by Saez and Zucman (2016), wealth returns should feature persistent heterogeneity (*type dependence*) and/or be an increasing function of wealth (*scale dependence*).

The model studied in this paper shares with the previous literature the presence of idiosyncratic returns to wealth. The main quantitative exercise is related to Benhabib et al. (2019) and Hubmer et al. (2020) who develop quantitative models to ascertain the importance of different factors to explain wealth concentration in the United States. Benhabib et al. (2019) consider an overlapping generations (OLG) economy with idiosyncratic earnings risk, non-homothetic preferences and idiosyncratic return risk across generations (but not within agents' life spans). As Benhabib et al. (2019) do not observe returns directly, the authors estimate the associated parameters by targeting U.S. wealth shares and data on social mobility. In this paper, I calibrate the parameters of the return process to match direct evidence on how returns vary along the US wealth distribution and do not target wealth shares per se.[7]

The main goal of Hubmer et al. (2020) is to examine changes in wealth inequality over time, but they also investigate the contribution of different channels for long-run inequality, including return heterogeneity. Compared to this paper, there are two main differences. First, Hubmer et al. (2020) calibrate a model of the U.S. economy using return estimates from Swedish data for the period 2000-2007. In this paper, I use U.S. return data to calibrate the return process. Second, I propose an alternative modeling of the individual return process that captures well the return heterogeneity observed in the data: it relies on *luck* and *type dependence*. The individuals who are "lucky" and/or have high return "types" save at a higher rate and end up at the top of the wealth distribution. In Hubmer et al. (2020), returns feature luck and *scale dependence*: being richer raises the expected return on wealth, which feeds back into wealth inequality. Conditional on the level of wealth, however, there are no persistent return differences between households.

The remainder of the chapter is organized as follows. Section 1.2 describes the data and the empirical evidence on returns to wealth in the United States. Section 1.3 outlines the model of household saving behavior. Section 1.4 describes the parameterization and the main quantitative results. Section 1.5 discusses alternative specifications of the return process and their distributional implications. Section 1.6 concludes.

[7]To be clear, I do not target any wealth share above the 50th percentile of wealth. I target the wealth share of the bottom 50% to pin down the borrowing constraint parameter.

1.2 Wealth and returns in the United States

I begin by describing the data and clarifying how wealth is defined. Then, I provide empirical evidence on how wealth is distributed in the United States and how the type of assets owned by households varies with their wealth. Finally, I construct a measure of wealth returns and investigate whether there is a systematic relationship between returns and households' position in the distribution of wealth.

1.2.1 Data sources and variable definitions

The primary data sources used in this study are the eleven waves of the Survey of Consumer Finances conducted every three years between 1989 and 2019. Each survey provides cross-sectional data on U.S. households' gross income for the calendar year preceding the survey, detailed information on their wealth and its components, as well as families' demographic characteristics. In addition to providing detailed information on household finances, one of the reasons the SCF is widely used to study wealth has to do with its sampling design. Because wealth is highly concentrated in the United States, the SCF oversamples wealthy households so that the collected data provides a good representation of the existing wealth and how it is distributed.[8]

The wealth concept used in this paper is marketable wealth, which is defined as the current value of all marketable assets less the current value of debts. I group wealth components into the following categories. Total assets are defined as the sum of (1) interest-earning assets[9]; (2) directly and indirectly held stocks (e.g., through mutual funds); (3) net equity in private businesses; (4) the value of real estate; (5) other miscellaneous financial assets; and (6) other nonfinancial assets[10]. Total liabilities are the sum of mortgage debt, consumer debt, including auto loans, and other debt such as educational loans ("debt").

1.2.2 Household wealth in the United States

To have an idea of the overall distribution of wealth in the United States, table 1.1 summarizes wealth shares in 2019 according to the SCF. The degree of concentration is striking: only 10% of the U.S. population own 76% of the total private wealth, while the poorest half of the population owns virtually no wealth (1.5%). Even within the richest 10%, wealth is unequally distributed: of the total wealth owned by the top 10%, 85% of it is concentrated in the hands of the richest 5%. Similarly, the richest 1% owns almost half of the wealth of the top 10%.

[8]See appendix 1.7.1 for a more detailed description of the SCF's sampling procedure.

[9]This category includes all types of transaction accounts (e.g., checking accounts, money market accounts, savings accounts), certificates of deposit, government bonds, corporate bonds, foreign bonds, other financial securities and the cash surrender value of life insurance plans.

[10]Vehicles are the main asset included in this category.

Table 1.1: Distribution of wealth in the United States, 2019 (SCF)

Percentile	Wealth share
Bottom 50%	1.5
Middle 40%	22.1
Top 10%	76.4
Top 5%	64.9
Top 1%	37.2

As previously mentioned, wealth is more concentrated than earnings. This is depicted in figure 1.1 which plots earnings shares and wealth shares in 2019 based on the SCF. Moreover, the large degree of wealth concentration is not specific to the year 2019. Although top shares increased somewhat between 1989 and 2019, the significant right skewness that characterizes the U.S. wealth distribution was already observed in the beginning of the sample period. Appendix 1.7.2 provides additional details of the evolution of U.S. wealth shares over time according to the SCF.

Figure 1.1: Wealth and earnings shares, 2019 (SCF)

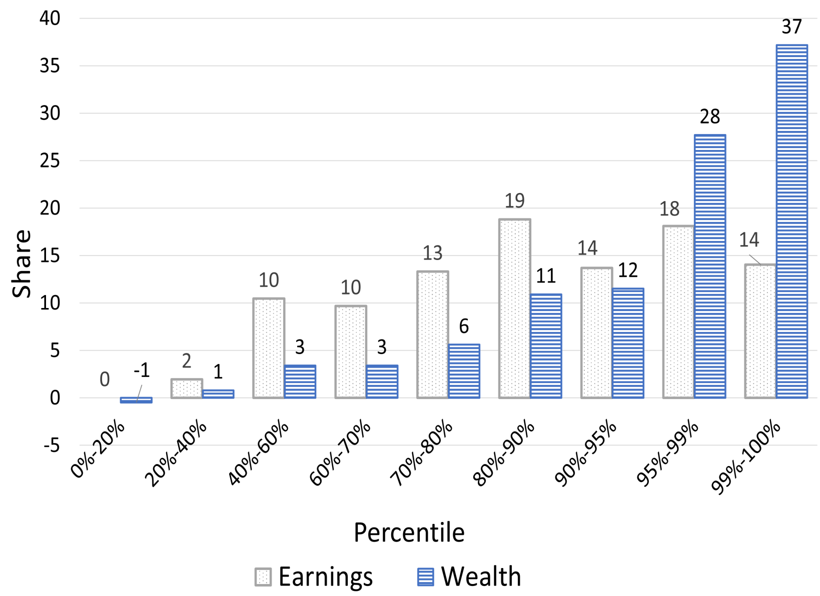

Household wealth is not just highly concentrated. Its composition is also very heterogeneous. This can be seen in figure 1.2, which plots the composition of gross asset portfolios for different percentiles of the wealth distribution. Real estate, especially primary homes, represent the main asset held by the bottom half of the distribution, who own very few public stocks and even less private equity. The other main assets owned are liquid assets and vehicles (included in "nonfinancial assets"). The middle class has more exposure to public stocks, but homes still dominate the asset portfolio, representing between 48% and 67% of the assets owned. Moving toward the top of the wealth distribution, the weight of housing shrinks gradually, while the importance of equity rises. This is especially true for private businesses which take up the

Figure 1.2: Gross asset portfolio composition by wealth percentile, 2019

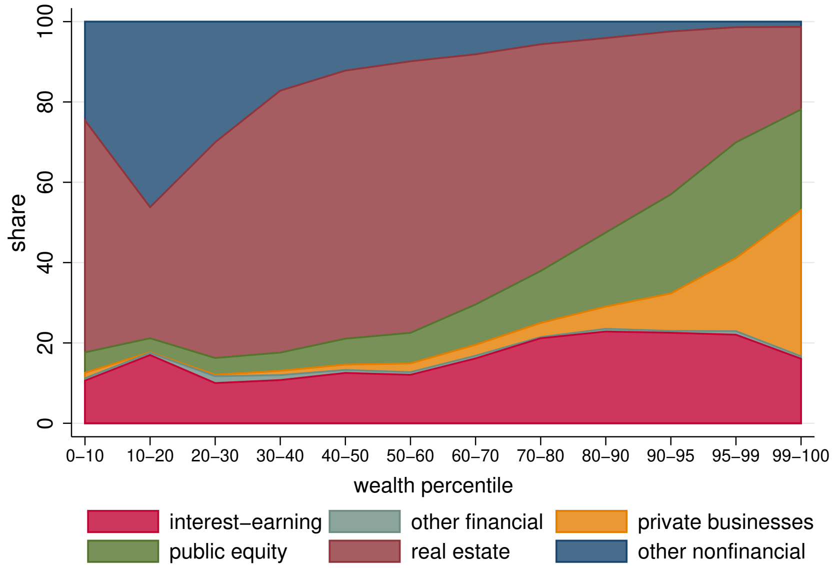

biggest share of the portfolio of the richest 1% (around 38%). Grouping public stocks and private businesses together implies that the richest 1% hold 64% of their assets in the form of equity.

1.2.3 Estimating returns to wealth

The goal for this section is to construct an estimate of the return on wealth and investigate the relationship between returns and wealth itself. To do this, I start by estimating the aggregate return on wealth and then compare it with the average return at different percentiles of the wealth distribution.

I start by clarifying how the return on wealth is defined. Denote by R_c the return on asset class c and ω_c its weight as a share of total wealth[11]. Then, the total return on wealth is defined as

$$R_w = \sum_c \omega_c R_c. \tag{1.1}$$

That is, the total return on wealth is given by the weighted average of the return on its different components.

In turn, the return on each wealth component is defined as the sum of two objects: (1) a yield component, capturing the net income generated by the asset; and (2) a capital gain component, reflecting fluctuations in its price. I use data from the SCF to estimate the yield component of each asset class (including debt), adapting the methodology of Moskowitz and

[11]The wealth components (or classes) considered here coincide with the ones described in section 1.2.1: interest-earning assets, public equity (stocks), private businesses, real estate, other financial and nonfinancial assets and debt. Debt enters with a negative sign.

Table 1.2: Yield component of returns, average 1990–2019

Wealth component	**Net income**	**Yield**
Interest-earning assets	Interest income	2.1%
Public equity	Dividends	1.8%
Private businesses	Net profits	9.0%
Real estate	Rental income	4.2%
Debt	Loan interest payments	2.7%

Vissing-Jørgensen (2002) and Kartashova (2014). Because this dataset does not include sufficient information to calculate unrealized capital gains or losses, I use the indices proposed by Shiller (2015) and data from the U.S. Financial Accounts[12] to estimate the capital gain component of real estate and equity assets. I now explain in further detail the computation of each of the two return components.

The yield component. Each wave of the survey provides information about the market value of each asset in the year of the survey and about the value of the associated income flow during the year preceding it. I use this data to calculate average annualized returns over three-year intervals which is the frequency of the data releases. To be clear, consider the following example using two consecutive waves of the SCF, 1989 and 1992. The average annualized return R over the period 1990-1992 is computed as the geometric average of returns R_1 and R_2 as follows:

$$R_1 = \left(1 + \frac{3NI_{1988}}{P_{1989}}\right)^{\frac{1}{3}} \tag{1.2}$$

$$R_2 = \left(1 + \frac{3NI_{1991}}{P_{1989}}\right)^{\frac{1}{3}} \tag{1.3}$$

$$R = (\sqrt{R_1 \cdot R_2} - 1) \cdot 100 \tag{1.4}$$

where NI denotes the total income flow generated by the asset and P represents the market value of the asset stock. Using equations (1.2)-(1.4), I construct an estimate of the return for each of the seven wealth components previously mentioned. Table 1.2 summarizes the income concept used in each asset category, as well as the resulting estimate of the average annualized yield return over the sample period. I assume that the categories "other financial assets" and "other nonfinancial assets" generate no income flows, which is why they are omitted from the table. See appendix 1.2 for a detailed description of the variables used to compute income flows and market values by asset category.

Private businesses were the highest-yielding asset, with an average yield return of 9.0% over the sample period. This is more than half of the yield return generated by real estate

[12]The U.S. Financial Accounts includes data on transactions and levels of financial assets and liabilities, by sector (e.g., households and nonprofit organizations and nonfinancial corporate businesses).

Table 1.3: Aggregate yearly return by wealth category, average over 1990-2019

Wealth component	**Yield**	**Capital gain**	**Return**
Interest-earning assets	2.1%	–	2.1%
Public equity	1.8%	4.9%	6.7%
Private businesses	9.0%	4.4%	13.4%
Real estate	4.2%	1.1%	5.3%
Debt	2.7%	–	2.7%
Other financial assets	–	0.4%	0.4%
Other nonfinancial assets	–	1.9%	1.9%

(4.2%) and about five times the yield on public equity (whose gains come mainly from price appreciations). The most relevant comparison, however, is of the total return on the different assets, for which we need to add the capital gain component. This is done as follows.

Capital gains and losses. To obtain the capital gain component of returns, I use the following sources. For public equity and real estate, I use the indices proposed by Shiller (2015)[13] and for private business equity, I use data from the U.S. Financial Accounts sponsored by the Federal Reserve Board.[14] I further assume no capital gains/losses on interest-earning assets and debt. Finally, I use the total value reported in the SCF to calculate capital gains or losses on other residual financial and nonfinancial assets.[15]

Over the past three decades, the largest capital gains were associated with public stocks and private businesses: an average gain of 4.91% and 4.39%, respectively. Real estate experienced a price boom during the first half of the 2000s, but the subsequent bust in 2008–10 explains the low overall gains during the whole period (1.10%). Finally, the residual categories of financial and nonfinancial assets earned an average capital gain of 0.39% and 1.87%, respectively.

The aggregate return on wealth. Having computed the yield and capital gain components, I obtain the total return on each asset by simply adding the two elements. Table 1.3 summarizes the average return of each asset category between 1990 and 2019, along with each of its components.[16]

On average, private businesses were the asset group that earned the highest returns, 13.4%, followed by publicly listed companies with a return of 6.7%. This premium is almost fully

[13] Up-to-date data from http://www.econ.yale.edu/∼shiller/data.htm.

[14] For noncorporate equity, I use the series "Nonfinancial noncorporate business; proprietors' equity in noncorporate business (wealth)" (NNBPEBA027N). The value of corporate equity is obtained from the series "Households and nonprofit organizations; corporate equities" (HNOCEAA027N). Both series are deflated by the CPI deflator.

[15] The overall results are similar if I assume no capital gains on other financial and nonfinancial assets. More generally, the results do not depend on the assumptions made regarding these assets because together they represent only about 6% of the total U.S. wealth.

[16] Appendix 1.7.4 provides further detail of the aggregate return on each wealth component, for each subperiod between 1990 to 2019. See table 1.16 and figure 1.12 in appendix 1.7.4 for a complete description of the estimates for all asset categories.

explained by differences in the yield return of the two assets (profits vs. dividends) and is in line with the findings of Kartashova (2014) for the period 1990–2010. Real estate in the United States earned an average return of 5.3% over the period, most of which is due to rents and not capital gains (which were lost in the housing market collapse of 2008–2010). Interest-earning assets yielded an average annual return of 2.1% which is, as expected, close to the estimate found for the average cost of debt, 2.7%. Finally, other financial and nonfinancial assets earned a yearly return of 0.4% and 1.9%, respectively.

The final step for estimating the aggregate return on wealth is to combine the return on each wealth category according to equation (1.1). Doing so for the aggregate U.S. economy yields an average yearly return of 6.80% between 1990 and 2019. Arriving at this point, one natural question emerges: have all U.S. households earned a return on their wealth close to 6.80%? If not, is there a relationship between returns and wealth?

1.2.4 Return heterogeneity along the wealth distribution

In this section, I investigate whether returns vary with wealth by repeating the calculations of the previous section at different percentiles of the wealth distribution. Specifically, I divide the population into nine wealth bins: 20%-40%, 40%-60%, 60%-70%, 70%-80%, 80%-90%, 90%-95%, 95%-97%, 97%-99% and 99%-100% percentiles.

The methodology is the same as previously. First, I estimate the yield component of returns for each wealth bracket applying equations (1.2)–(1.4). The second step is to impute the capital gains estimated in section 1.2.3 to each asset category. Finally, the total wealth return is obtained by weighting the return on the different wealth components according to equation (1.1). The results are depicted in figure 1.3.

Figure 1.3: Average return on wealth by percentile of wealth

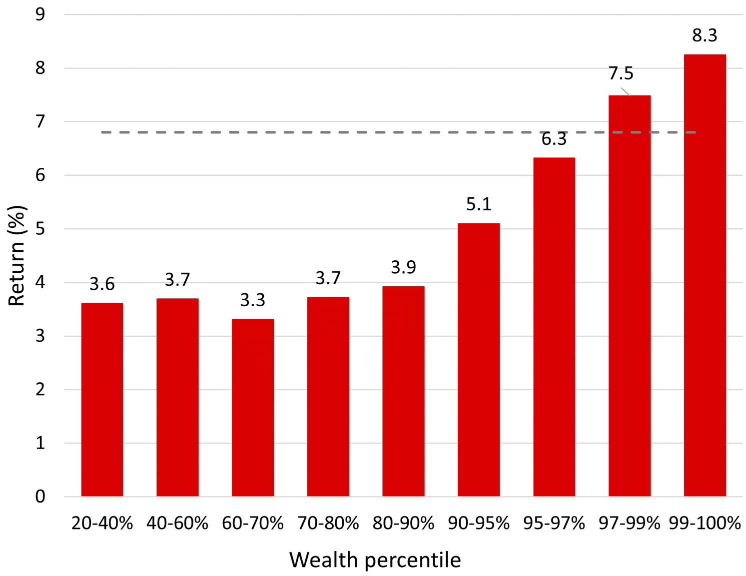

Clearly, not all households earn the same return on wealth. Figure 1.3 depicts a clear positive correlation between average returns and the households' wealth percentile. Between 1990 and 2019, households in the 20%-40% percentiles of the wealth distribution received an average return of 3.61%, which is almost half of the aggregate return represented by the grey dashed line. In contrast, the richest 1% earned an average return of 8.25%, implying a difference of 4.64 percentage points with respect to the bottom group. For comparison, in Norway (2005–2015), Fagereng et al. (2020) estimate that the average return on wealth rises from -1.5% for the 20th percentile to 3.8% for the 50th percentile of the wealth distribution, and it further rises to 5.7%, approximately, for the 99th percentile.[17] Using Swedish data from 2000-2007, Bach et al. (2020) estimate that the expected excess return on wealth rises from 3.8% for the 20%–30% wealth percentile to 4.7% for the 50%-60% wealth percentile. It then rises further and varies between 6.58% and 8.32% for individuals within the top 1%.[18]

Next, I discuss how these return differentials can partially be explained by the heterogeneous portfolio composition documented in section 1.2.2. However, this is not the only source of heterogeneity.

1.2.5 Sources of return differentials

In this section, I decompose the return differentials displayed in figure 1.3 into two sources: (1) heterogeneous composition of wealth portfolios and (2) return differences within asset categories.

Starting with the composition of wealth portfolios, section 1.2.2 has shown that there are systematic differences between individuals. High-wealth households own relatively more equity, while low and middle-wealth households are more exposed to real estate assets. Given that equity earns, on average, higher returns than real estate, it is unsurprising that wealthier households earn larger wealth returns.

In addition, richer households earn higher returns on some assets. This is true for private businesses, real estate and even interest-earning assets, although the premium is smaller. This is depicted in figure 1.4 which plots the average return on selected asset categories for different wealth percentiles. The largest differentials are observed on private businesses, where households in the top 10% of wealth earn higher returns than the next 20% (up to 6.9 percentage points more)[19]. Curiously, the relationship between returns and wealth is not monotone at the top: households within the top 1% earn lower returns than the next 9%, a difference that goes up to 2.8 percentage points. The same broad pattern is observed in real estate. Moving from the

[17] Fagereng et al. (2020) use a slightly different definition of returns in which income is divided by gross wealth, instead of net wealth.

[18] Excess returns are defined relative to the Swedish Treasury bill which earned, on average, 1.5% per year over the period.

[19] The sample of private equity owners within the wealth percentile groups 20%–40%, 40%–60% and 60%–70% is substantially smaller than that of the other wealth groups and leads to very volatile return estimates. Therefore, I do not consider them in the calculation of private businesses' returns.

Figure 1.4: Returns by asset class and percentile of wealth

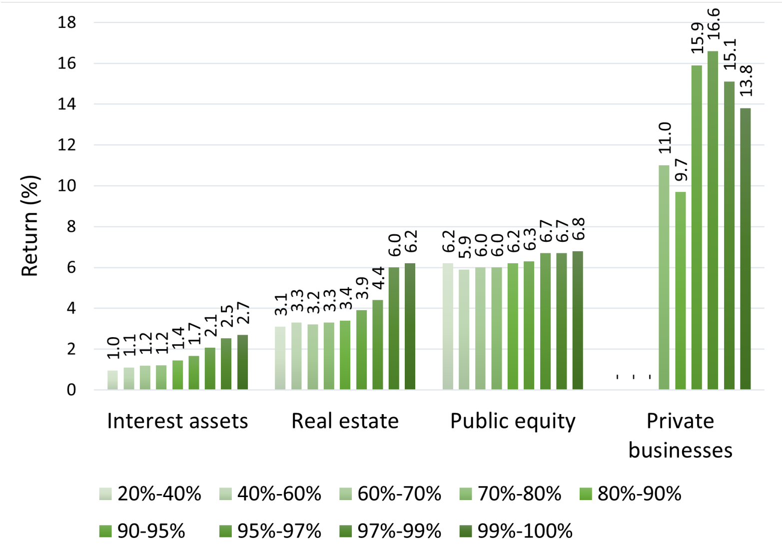

bottom 20% of the wealth distribution to the top 1% roughly doubles housing returns, from 3% to 6.2%, suggesting that wealthy households are able to extract relatively more income from their properties. Even on interest-earning assets there is a "wealth premium" that goes up to 1.7 percentage points between the lowest and highest wealth groups considered. In contrast, return differentials are much lower for investments in public stocks, supporting the idea that stock portfolios are better diversified for all households.

What is the contribution of each of these factors to the final return differentials? One way of ascertaining the contribution of portfolio composition is to shut down all portfolio differences among households and calculate the resulting (counterfactual) return on wealth. Similarly, one could shut down all heterogeneity in returns within asset classes, while still allowing for differences in portfolio composition and build a counterfactual return estimate. These counterfactual returns are represented by the blue and green bars in figure 1.5, which are contrasted with the actual returns of each wealth group. I also include the aggregate return on wealth, 6.8%, for reference (dashed line).

To build the counterfactual estimates represented by the green bars, I allow households to earn different returns within asset classes but impose that all households own the aggregate wealth portfolio.[20] For households up to the 97% percentile, imposing the aggregate portfolio implies raising the relative importance of equity (to the detriment of real estate) and reducing debt exposure, which contributes to raise the average return by up to 1.7 percentage points. The wealthiest 1% households would be the most hurt by the imposition of the aggregate portfolio, as their average return would fall from 8.4% to 7.3%. This is mainly due to the substitution of private businesses with lower yielding assets.

If, instead, we eliminate return differences on similar assets, we obtain the estimates repre-

[20]I assume that if households in the 20%–70% wealth percentiles owned private businesses, they would earn the same return as the next wealth bracket (70%-80%).

Figure 1.5: Wealth returns by percentile of wealth: actual vs. counterfactuals

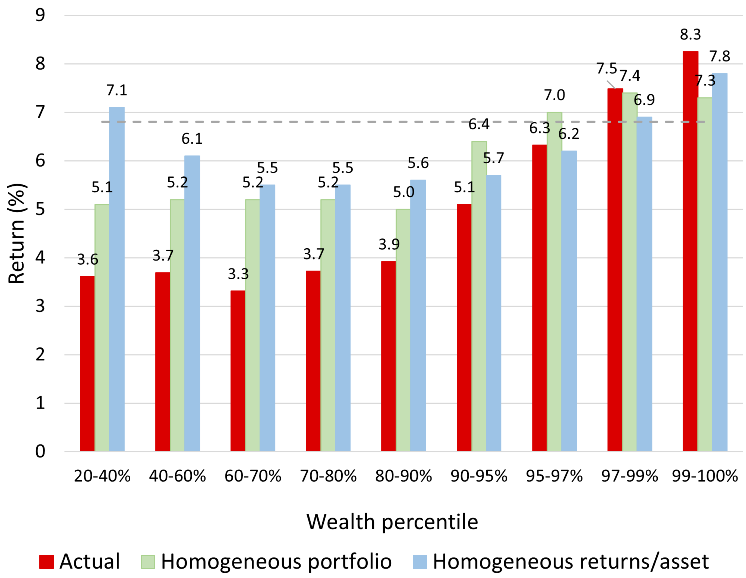

sented by the blue bars. Bottom households, who earn lower returns on all assets, would gain the most: a rise of 3.4 percentage points in the average return. The main effect comes from the resulting increase in housing returns at the bottom. Similarly, the middle class would see their wealth return increase. For example, the 40%-60% wealth percentiles would see the average return rise from 3.8% to 6.1%. In contrast, the wealthiest 5% households would see their returns fall up to 0.7 percentage points, after losing the return advantages displayed in figure 1.4.

Overall, both factors are quantitatively important to explain wealth return differences. Return heterogeneity within asset classes is particularly important for households at the bottom of the distribution (the 20%–60% percentiles). Eliminating return differentials on similar assets would break the positive correlation between returns and wealth and would bring to 0.7 percentage points the return gap between the bottom 20%–40% and the top 1% of wealth. However, some important return differences would remain: families in the middle of the wealth distribution would earn a 5.5% annual return on wealth, which is 2.3 percentage points lower than the return that the richest 1% would earn in this counterfactual scenario. The remaining differential is explained by the heterogeneous portfolio composition. If we were to shut down all heterogeneity in the composition of wealth portfolios instead, we would still observe a positive correlation between returns and wealth, but the return differential between bottom and top households would be reduced from 4.7 to 2.2 percentage points. Thus, it is important to consider both sources of return differentials to understand the magnitude of the overall return differences.

1.2.6 The correlation between returns and wealth

In the analysis of the previous section, I uncovered important differences in the rate of return that U.S. households receive on their investments. In particular, the analysis showed that richer households earn, on average, a higher return on their wealth. I now discuss three mechanisms, proposed by Gabaix et al. (2016), that generate a positive correlation between wealth and returns: (1) *type dependence*, (2) *scale dependence* and (3) *luck*.

Type dependence refers to the possibility that households face fundamentally different return processes. Some individuals may be "high return types," for example, reflecting their education, talent as investors or risk tolerance. "High-type" individuals earn higher average returns than "low types", which allows the high types to accumulate relatively more wealth over time.[21] Therefore, this mechanism can rationalize the empirical correlation between average returns to wealth and the households' position in the wealth distribution: "high-types" earn persistently higher returns, accumulate more wealth and end up at the top of the wealth distribution.

Alternatively, returns may feature *scale dependence*. As the term suggests, in this case the *level* of wealth matters for the return on wealth.[22] This mechanism can rationalize that the return of an investment depends on its size or that some instruments have minimum investment requirements (or other types of "barriers" to access). Wealthier individuals earn higher returns precisely because they have more wealth. In turn, higher returns allow individuals to accumulate relatively more wealth which raises their expected return and so on.

Finally, the positive correlation between realized returns and wealth may simply be a product of *luck*, that is, idiosyncratic randomness. Suppose that returns are stochastic and idiosyncratic. Even if the return process is identical for all individuals (ex-ante), the ex-post realization of returns varies with the realization of the idiosyncratic shock. Those households who are "lucky" and receive higher return realizations can accumulate relatively more wealth. Once again, this would generate a positive correlation between realized returns and wealth.

Without the panel dimension in the SCF, it is challenging to disentangle the contribution of the previous mechanisms for the return heterogeneity measured in the data. Nonetheless, the empirical evidence of Bach et al. (2020) and Fagereng et al. (2020) suggests that all three mechanisms are likely to affect returns in practice. Bach et al. (2020) investigate the presence of type and scale effects by comparing the expected return on wealth of pairs of twins. The authors' strategy relies on the assumption that twins share the same investment "type," and then they estimate the relationship between the expected returns and scale.[23] Bach et al. (2020) find that scale dependence, type dependence and transitory variation explain 7%, 16% and 77%,

[21] Different "types" may differ, not only in the expected return but also in the standard deviation of the return innovations.

[22] Again, scale may affect the mean and the standard deviation of individual returns.

[23] As recognized by the authors, this strategy is likely to underestimate the contribution of household type dependence if twins do not fully share the same investment *type*.

respectively, of the variance in wealth returns. The estimation of Fagereng et al. (2020) suggests that both scale and individual fixed effects are statistically significant. For example, the results imply that type dependence and transitory variation explain 52% of the 20 percentage point return difference between the 15% and 85% wealth percentiles. Scale dependence accounts for the remaining difference. All in all, the existing empirical evidence does not provide a precise estimate of the relative importance of type dependence, scale dependence and luck. However, the evidence suggests that all mechanisms are realistic features of returns to wealth in the data.

Heterogeneous returns and wealth inequality. How does return heterogeneity affect the distribution of wealth in the United States? In the remainder of the paper, I address this question through the lens of a model of household wealth accumulation. As previously discussed, there are different mechanisms that can replicate the empirical positive correlation between returns and wealth. I propose a model in which returns feature *type dependence* and *luck*. This choice is motivated by the following observations. First, I find that *luck* alone is insufficient to accurately replicate the return differences estimated in the previous section (this is discussed in section 1.5). Second, I find that a simple formulation with only few different return types is able to match the empirical patterns of wealth returns in U.S. data[24]. Finally, the exogenous return types seem like a natural first step to capture the large return differences observed in the data, given the limited information on the relative importance of "type" and "scale" dependence. Notice that the main difference between the two mechanisms is that, while return "types" induce fully exogenous return differences, "scale" dependence induces a strategic behavior of individuals who internalize that their saving behavior affects the expected return on wealth (and potentially return volatility). This implies that the specific way in which returns depend on scale matters for the optimal saving behavior and, thus, for the distribution of wealth. Given the lack of empirical evidence on the strength of such strategic behavior, I choose to model persistent return differences as heterogeneous agent "types".

1.3 A model of wealth inequality and heterogeneous returns

The basic building block is the workhorse incomplete-market economy with idiosyncratic labor income risk. To account for the return differentials observed in the data, I amend the workhorse model to feature return heterogeneity.

1.3.1 Setup

Individuals. The economy is populated by a continuum of individuals indexed by i who choose the path of consumption that maximizes

[24] Again, this is just one way of generating persistent return differences. It does not imply that scale dependence is not empirically relevant. As discussed in the main text, both types of mechanisms — types and scale — are likely to play a role in practice.

$$\mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_{it}) dt, \tag{1.5}$$

where $c_{it} \geq 0$ is consumption, and ρ is the discount rate. Time t is continuous, and preferences display constant relative risk aversion (CRRA); that is, $u(c) = \dfrac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$.

Individuals accumulate wealth a_{it} over time according to

$$\dot{a}_{it} = y_{it} + r_{it} a_{it} - c_{it} \tag{1.6}$$

where y_{it} denotes labor income, and r_{it} is the return on wealth. Moreover, individuals face a borrowing limit

$$a_{it} \geq \underline{a}, \tag{1.7}$$

with $-\infty < \underline{a} \leq 0$.

Labor income y_{it} is idiosyncratic and exogenous. It evolves stochastically over time according to the stationary diffusion process

$$dy_{it} = \mu_y(y_{it}) + \sigma_y(y_{it}) dW_{it}, \tag{1.8}$$

where W_{it} is a standard Brownian motion.[25] The functions μ_y and σ_y respectively represent the drift and the diffusion of the process. Section 1.4 makes specific assumptions on their functional forms, which will determine the mean and standard deviation of the growth rate of earnings.

Likewise, the return r_{it} is idiosyncratic, exogenous and stochastic. Relative to earnings, I consider a more flexible formulation that allows the drift and the diffusion of the return process to potentially differ across individuals:

$$dr_{it} = \mu_{r,i}(r_{it}) + \sigma_{r,i}(r_{it}) dZ_{it}, \tag{1.9}$$

where Z_{it} is a standard Brownian motion. This formulation, known as *type dependence*, allows for the presence of "high" and "low" return *types*.[26] If there is only one *type* (i.e., $\mu_{r,i} = \mu_r$, $\sigma_{r,i} = \sigma_r \ \forall i$), then all individuals face the same return process, and all return heterogeneity is due to the ex-post realization of the shock. If there are more return *types*, then there is also some ex-ante heterogeneity in the sense that individuals face different return processes.

Individuals maximize equation (1.5) subject to equations (1.6) and (1.7) and the exogenous processes for y_i and r_i described by equation (1.8) and equation (1.9). The state of the economy

[25] A standard Brownian motion (or Wiener process) is a stochastic process W that satisfies $W(t+\Delta t) - W(t) = \epsilon\sqrt{\Delta t}$, where $\epsilon \sim \mathcal{N}(0,1)$.

[26] Gabaix et al. (2016) were the first to propose such formulation of returns, arguing that it can help explain fast changes in top inequality as suggested by empirical evidence. Fagereng et al. (2020) and Bach et al. (2020) support this formulation of returns by showing that return differences in Nordic countries are highly persistent across the entire support of the wealth distribution.

is the joint distribution of wealth, income and returns.

1.3.2 Stationary equilibrium

For convenience, I now suppress the subscript i on wealth, income and returns and keep it only to identify individuals' return *type*. Individuals' consumption-saving decisions and the probability density function over wealth, income and returns can be summarized with a Hamilton-Jacobi-Bellman (HJB) equation and a Kolmogorov Forward (KF) equation[27]:

$$\begin{aligned}\rho v_i(a,y,r) =& \max_c u(c) + \partial_a v_i(a,y,r)(y+ra-c) + \partial_y v_i(a,y,r)\mu_y(y) \\ &+ \partial_r v_i(a,y,r)\mu_{r,i}(r) + \frac{1}{2}\partial_{yy} v_i(a,y,r)\sigma_y^2(y) + \frac{1}{2}\partial_{rr} v_i(a,y,r)\sigma_{r,i}^2(r) \qquad ,\forall i\end{aligned} \tag{1.10}$$

$$\begin{aligned}0 =& -\partial_a[s_i(a,y,r)g_i(a,y,r)] - \partial_y[\mu_y(y)g_i(a,y,r)] - \partial_r[\mu_{r,i}(r)g_i(a,y,r)] \\ &+ \frac{1}{2}\partial_{yy}[\sigma_y^2(y)g_i(a,y,r)] + \frac{1}{2}\partial_{rr}[\sigma_{r,i}^2(r)g_i(a,y,r)] \qquad ,\forall i\end{aligned} \tag{1.11}$$

where $s_i(a,y,r) \equiv y + ra - c_i(a,y,r)$ is the saving policy function, and optimal consumption satisfies $c_i(a,y,r) = (u')^{-1}(\partial_a v_i(a,y,r))$.

Intuitively, the HJB equation (1.10) characterizes the optimal consumption and saving behavior of individuals, given the exogenous processes for earnings and returns. The KF equation (1.11) describes the evolution over time of the probability density function $g_{it}(a,y,r)$, given optimal individual saving decisions and the evolution of earnings and returns. In the stationary equilibrium, $\partial_t g_{it}(a,y,r) = 0$ which explains the left-hand side of equation (1.11).

1.4 Quantitative analysis

The objective of this section is to quantify the importance of return heterogeneity for the distribution of wealth in the United States. I first discuss the parameterization strategy of the model presented in section 1.3 and then discuss the extent to which it is able to replicate a set of moments from U.S. data.

1.4.1 Model parameterization

The parameterization proceeds in two steps. First, I calibrate a set of parameters outside the model using estimates from the literature. Then, the remaining parameters are jointly calibrated to match several moments in the data.

[27] See Achdou et al. (ming) for an intuitive derivation of the HJB and KF equations in a slightly simpler setting.

Externally calibrated parameters

The model is calibrated at an annual frequency. The parameters of the CRRA utility function and the earnings process closely follow the existing literature. The coefficient of relative risk aversion, γ, is set to 2. The earnings process is based on the traditional log-normal framework. That is, I assume that log-earnings, z_t, follow an Ornstein-Uhlenbeck (OU) process[28]:

$$dz_t = \theta_z(\bar{z} - z_t)dt + \sigma_z dW_t. \tag{1.12}$$

I set parameter θ_z equal to 0.11 to match an autocorrelation of log-earnings equal to 0.9, and the standard deviation of innovations σ_z is equal to 0.2[29]. Finally, $\bar{z}$ is set to 0.78 to ensure that the aggregate earnings sum up to 1 (normalization).

Fitted parameters

The remaining parameters are the discount rate ρ, the borrowing constraint $\underline{a}$ and the parameters of the return process.

Discount rate. I target an aggregate rate of return of 6.80% to pin down the discount rate ρ. This is the estimate found in section 1.2.3 for the United States between 1990 and 2019.

Borrowing constraint. The borrowing constraint $\underline{a}$ is chosen to match the share of wealth of the bottom 50% in 2019 according to the SCF.

Return process. I assume that returns follow an Ornstein-Uhlenbeck process and partition the population into three return *types* indexed by j of size δ_j each ($J = 3$).[30] Thus, an individual of type j faces the return process

$$dr_t = \theta_r(\bar{r}_j - r_t)dt + \sigma_{r,j} dZ_t. \tag{1.13}$$

There are nine return parameters to be estimated: θ_r, $\bar{r}_j$, $\sigma_{r,j}$, for $j = 1, 2, 3$ and δ_1, δ_2.[31] To estimate these parameters, I target the empirical average return of nine different wealth brackets. More specifically, I target the average return of the 20%-40%, 40%-60%, 60%-70%, 70%-80%, 80%-90%, 90%-95%, 95%-97%, 97%-99% and top 1% wealth percentiles. I consider narrower wealth brackets at the top of the distribution to account for the fact that empirical returns exhibit more heterogeneity at the top of the wealth distribution as depicted in figure 1.3.

[28] The OU process is the continuous-time analogue of a discrete-time AR(1) process.
[29] As in Aiyagari (1994) and Guvenen et al. (2019), for example.
[30] In section 1.5, I discuss the choice of the number of return types and the implications of alternative choices.
[31] As the population size is normalized to 1, the mass of type 3 individuals is given by $1 - \delta_1 - \delta_2$.

Overall, I target 11 moments and estimate 11 parameters. The optimization routine is as follows. For arbitrary values of the vector of parameters $\Theta = (\rho, \underline{a}, \theta_r, \bar{r}_1, \bar{r}_2, \bar{r}_3, \sigma_{r,1}, \sigma_{r,2}, \sigma_{r,3}, \delta_1, \delta_2)$, the model is solved using the algorithm described in appendix 1.8.1. This yields the optimal individual decision rules and the stationary distribution over wealth, earnings and returns. Using these objects, I compute the model implied aggregate rate of return, the wealth share of the bottom 50% and the average return of the nine wealth percentiles previously specified. Then, the fitted parameters $\hat{\Theta}$ are the ones that minimize the distance between the model-generated moments and the targeted moments from U.S. data. Formally, let $M(\Theta)$ denote the vector of empirical moments targeted in the calibration and let $\hat{M}$ denote the corresponding vector of moments generated by the model. Then,

$$\hat{\Theta} = arg\min_{\Theta}(\hat{M} - M(\Theta))'W(\hat{M} - M(\Theta)), \tag{1.14}$$

where W is the weighting matrix which I set equal to the identity matrix, $W = I$. Additional details of the optimization routine are presented in appendix 1.8.2.

Table 1.4 presents an overview of the externally calibrated and internally fitted parameters. Before discussing them in more detail, the next section shows how well the model replicates the empirical moments targeted in the calibration.

Table 1.4: Overview of parameters

Parameter	Description	Method	Value
γ	Coefficient of relative risk aversion	External	2
θ_z	Persistence of log-earnings	External	0.11
$\bar{z}$	Aggregate earnings	External	0.78
σ_z	Standard deviation of earnings innovations	External	0.2
ρ	Discount rate	Internal	0.088
$\underline{a}$	Borrowing constraint	Internal	-0.65
θ_r	Persistence of returns	Internal	3.08
$\bar{r}_1$	Mean return type 1	Internal	0.033
$\bar{r}_2$	Mean return type 2	Internal	0.058
$\bar{r}_3$	Mean return type 3	Internal	0.082
$\sigma_{r,1}$	Diffusion return type 1	Internal	0.056
$\sigma_{r,2}$	Diffusion return type 2	Internal	0.202
$\sigma_{r,3}$	Diffusion return type 3	Internal	0.057
δ_1	Mass of type 1 agents	Internal	0.80
δ_2	Mass of type 2 agents	Internal	0.18

1.4.2 Model fit

The estimated baseline model captures the targeted moments quite well. Table 1.5 shows that the model matches the targeted aggregate rate of return and the share of wealth of the bottom 50%. Figure 1.6 shows, graphically, that the fit is also very good for the remaining return moments targeted in the calibration.

Table 1.5: Targeted moments I

	Model	Data
Aggregate return	6.79%	6.80%
Wealth bottom 50%	1.5%	1.5%

Figure 1.6: Targeted moments II: average return by wealth percentile

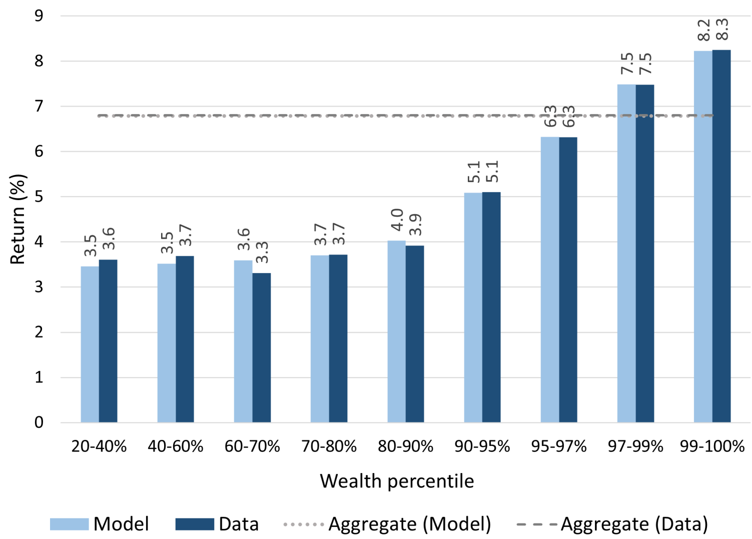

1.4.3 Estimated return heterogeneity

For convenience, table 1.6 reproduces the fitted return parameters. These are consistent with the majority of U.S. households, 80%, being of the "low" type (type 1). These households have "low" expected returns (an unconditional mean of $\bar{r}_1 = 0.033$) with low volatility ($\sigma_{r,1} = 0.056$). The value of θ_r determines the strength with which the process reverts to its mean and thus, is closely related to the persistence of the process. The value $\theta_r = 3.08$ implies an autocorrelation of about 0.05 which, in turn, implies that return shocks are very transitory. Type 2 households, or "middle" types, represent about 18% of the population. Their return process is associated with a higher mean ($\bar{r}_2 = 0.058$) and higher volatility ($\sigma_{r,1} = 0.202$) of returns than "low" type households. By assumption, the persistence of return shocks is identical for

Table 1.6: Overview of return parameters

	Low (type 1)	Middle (type 2)	High (type 3)
Mean, $\bar{r}_j$	0.033	0.058	0.082
Standard deviation, $\sigma_{r,j}$	0.056	0.202	0.057
Persistence, θ_r	3.08	3.08	3.08
Share, δ_j	0.80	0.18	0.02

all types, and therefore, it is also low for the "middle" types. Finally, "high" types represent only 2% of the population and have "high" expected returns ($\bar{r}_3 = 0.082$) with relatively low volatility ($\sigma_{r,3} = 0.057$).

The estimated parameters have no direct counterpart in the data to which they can be compared. Nonetheless, I briefly comment on how some return statistics implied by the model compare to the available evidence from the Nordic countries. For Sweden (2000-2007), Bach et al. (2020) find that the expected return on wealth rises from 5.3% for the 20%-30% wealth percentiles, to 7.5% to the 97.5%-99% wealth percentiles. It then rises further for individuals within the top 1% of the wealth distribution, fluctuating between 8.1% and 9.8%. For Norway (2005–2015), Fagereng et al. (2020) estimate that the average return on wealth rises from approximately -1.8% in the 20% wealth percentile, to 5.7% at the top 1% wealth percentile[32]. In my model, the average return on wealth rises from 3.4% in the 20%-40% wealth percentiles, to 8.2% for the top 1% of the wealth distribution.

Table 1.7: Idiosyncratic volatility of returns

Wealth percentile	20%	90%	99%
Model	6.5%	14.5%	5.8%
Bach et al. (2020)	8.0%	6.0%	8.7%-27.5%

Turning to the relationship between return risk and wealth, Bach et al. (2020) provide an estimate of the idiosyncratic volatility of returns and how it varies with wealth. I compare this object with the volatility of returns implied by σ_r (averaged over the different types) in the model. Table 1.7 depicts these statistics. The model implies a relatively larger idiosyncratic volatility of returns at the 90th percentile of the wealth distribution, and a lower volatility of returns for individuals in the top 1%[33]. Overall, the model is broadly consistent with the evidence from Nordic countries pointing towards a positive correlation between wealth, on one

[32] Based on figure 5 in Fagereng et al. (2020)

[33] In appendix, I provide additional evidence on the cross-sectional standard deviation of returns implied by the model and how it compares with the empirical evidence of Bach et al. (2020) and Fagereng et al. (2020).

hand, and the average return and the idiosyncratic volatility, on the other hand. The exception is at the top 1% of the wealth distribution, where the model implies considerably less idiosyncratic return risk.

It is important to recognize that the (time-series) return process faced by individuals in the model is calibrated to match empirical moments from cross-sectional data. This is not ideal as it would be preferable to use panel data to identify the different components of the return process. In particular, the cross-sectional moments cannot disentangle between ex-ante return differences, captured by return types, and the ex-post return differences generated by the stochastic return component. To address this issue, I consider alternative combinations of ex-ante and ex-post heterogeneity by allowing for different numbers of return "types". In section 1.5, I show that, while different specifications lead to different parametrizations of the return process (and relative importance of ex-ante versus ex-post differences), the implications for the overall distribution of wealth are broadly similar[34]. I start by presenting the baseline model's implications for long-run wealth inequality in the next section.

1.4.4 Results: steady-state wealth inequality

Figure 1.7 plots the Lorenz curve generated by the model (solid line) and, for comparison, its empirical counterpart in 2019 (dashed line). For further detail, table 1.8 summarizes the wealth shares of selected groups.

Figure 1.7: Lorenz curve: model and data

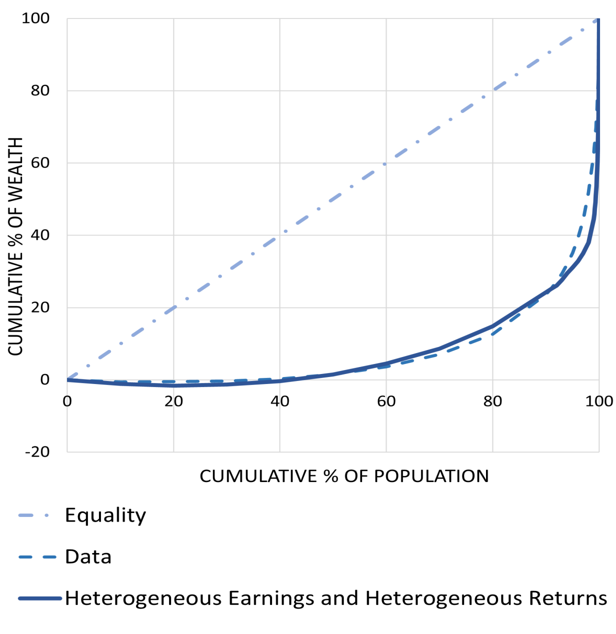

The degree of wealth concentration implied by the baseline model is not only large but

[34]Conditional on matching the targeted cross-sectional return moments.

Table 1.8: Wealth shares: model and data (2019)

Percentile	Model	Data
Bottom 50%	1.5	1.5
Middle 40%	22.8	22.1
Top 10%	75.7	76.4
Top 5%	68.9	64.9
Top 1%	55.5	37.2

also remarkably close to the empirical wealth shares in the data. This is true although empirical wealth shares were not directly targeted in the calibration, with the exception of the bottom 50% as a whole[35]. In the model, the wealthiest 10% households own 75.7% of the total wealth. This value is almost identical to the empirical top 10% share observed in 2019: 76.4%. Between the 50th and 90th wealth percentiles, the model tracks the distribution of wealth observed in the data very closely as is clear in figure 1.7. Moving further up the distribution, the model-implied wealth share of the richest 5% is also very close to its empirical counterpart (68.9% vs. 64.9%). For the richest 1%, the model implies an even greater wealth share than what is observed in the data (55.5% vs. 37.2%), implying even more wealth concentration *within* the top 5% than what is measured in the data.

To understand the role of *type dependence* for wealth inequality, the distribution of return types over different wealth percentiles is plotted in figure 1.8. The main observation is that the household's return type is positively correlated with the household's position in the distribution of wealth. "Low" types are more likely to belong to the bottom or the middle of the wealth distribution. This is clear from the first two blue bars on the left which indicate that 53% and 44% of low-type households end up in the bottom 50% and the middle 40% of the wealth distribution, respectively. In contrast, high-type individuals are extremely likely to belong to the top 10% of the wealth distribution. In fact, 99% of high types will end up within the top 5% of the wealth distribution, and 59% of them will belong to the top 1%. Finally, middle-type families are likely to become middle class households, although there are middle types at the bottom 50% and at the top 10% or the top 5%.

Overall, the main takeaway from the baseline model is that adding return heterogeneity, consistent with the U.S. data, to the workhorse model of earnings inequality generates top wealth shares that are remarkably close to their empirical counterparts. In the next section, I investigate just how important return differences are to understand wealth inequality in the US.

[35] Appendix 1.8.4 shows an alternative calibration in which I do not target the bottom 50% empirical share, and assume that households cannot borrow, that is, $\underline{a} = 0$. The wealth distribution of the bottom 50% is somewhat larger than in the baseline model (6.8%), mainly at the expense of the next 40% (19.4% versus 22.1% in the baseline). The model-implied top 10% share is equal to 73.8%.

Figure 1.8: Distribution of return types over different wealth percentiles

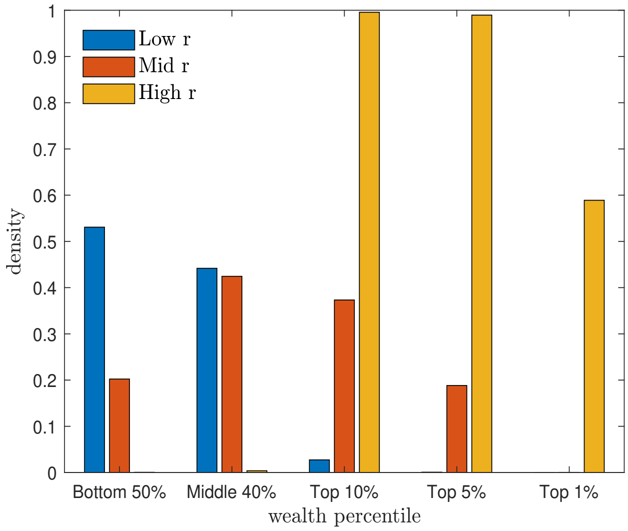

1.4.5 Counterfactual: homogeneous returns

In this section, I present a counterfactual exercise whose goal is to understand the relative importance of heterogeneous returns for long-run wealth inequality. Specifically, I shut down all sources of return heterogeneity ($\mu_{r,i} = \mu_r$ and $\sigma_{r,i} = 0\ \forall i$) and re-estimate the model. I set $\rho = 0.088$ and $\mu_r = 6.79\%$ to match the aggregate rate of return in the baseline economy. Then, I re-calibrate $\underline{a}$ to match the bottom 50% wealth share which yields $\underline{a} = -3.99$. Table 1.9 summarizes the return parameters associated with this counterfactual exercise.

Table 1.9: Overview of return parameters: homogeneous returns

	Low (type 1)	**Middle** (type 2)	**High** (type 3)
Mean, $\bar{r}_j$	0.0679	0.0679	0.0679
Standard deviation, $\sigma_{r,j}$	0	0	0
Persistence, θ_r	-	-	-
Share, δ_j	0.80	0.18	0.02

The implications of ignoring return heterogeneity are depicted in figure 1.9 and table 1.10. Figure 1.9 plots the Lorenz curve associated with the counterfactual economy with earnings inequality but no return heterogeneity (dotted line). For comparison, it also displays the corresponding object for the baseline model and the data. As the bottom 50% share is a target in the calibration, the relevant comparison is that of wealth shares *within* the top 50% of the distribution. Shutting down return heterogeneity has large implications for the distribution of wealth which becomes much less concentrated. For example, the wealth share of the middle

Figure 1.9: Distribution of wealth: data and model under different assumptions

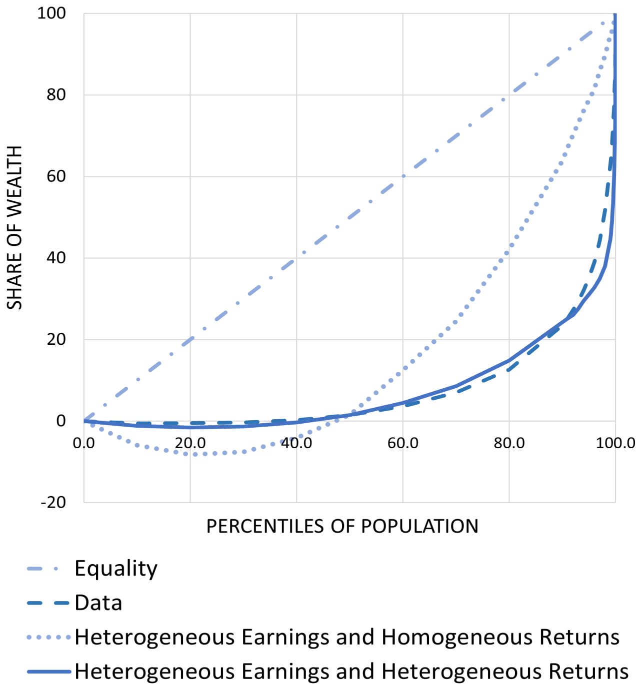

Table 1.10: Wealth shares: homogeneous returns, baseline and data

	Homogeneous returns	Baseline	Data
Bottom 50%	1.5	1.5	1.5
Middle 40%	62.3	22.8	22.1
Top 10%	36.2	75.7	76.4
Top 5%	21.1	68.9	64.9
Top 1%	5.2	55.5	37.2

40% increases from 22.8% in the baseline to 62.3% in the homogeneous return specification. In contrast, the share of the richest 10% drops roughly by half, from 75.7% to only 36.2%. Moving further up the distribution, the relative drop is progressively larger, indicating that return heterogeneity is of first-order importance to understand top wealth shares. Specifically, the top 5% owns only 21.1% of the total wealth under homogeneous returns, which is far from the 64.9% observed in the data and the 68.9% implied by the baseline model. Similarly, the top 1% is predicted to own 5.2% of the total wealth when the corresponding empirical and baseline shares are 37.2% and 55.5%.

The main conclusions from this exercise are summarized as follows. First, labor income differences are insufficient to explain large top wealth shares, in line with the general findings of the literature. Second, this is increasingly true as one moves further up the wealth distribution and earnings risk becomes weaker as a source of wealth dispersion. Finally, considering return heterogeneity on top of earnings inequality can fully explain the degree of wealth concentration observed in the data, suggesting that both factors must be taken into account in order to understand long-run wealth inequality.

1.5 Discussion: cross-sectional return heterogeneity and type dependence

In the baseline model, returns feature *type dependence* and I assume that there are three return *types* in the population. This modeling choice is justified by the following observations.

Figure 1.10: Model fit (one type): average return by wealth percentile

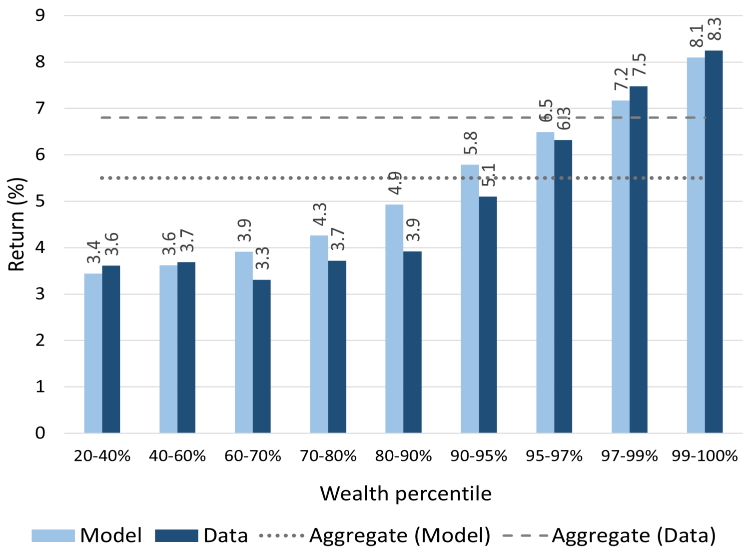

First, the model with only one return type (i.e. no type dependence) performs poorly on the matching of the targeted empirical moments. This is depicted in figure 1.10.[36] The matching is particularly poor for the targeted aggregate rate of return (5.5% vs. 6.8%) and the average return of percentiles 60%-70%, 70%-80%, 80%-90% and 90%-95%, which are too high in the model compared to the data. This result is perhaps not surprising. In this specification, there are only three return parameters: $\bar{r}$, σ_r and θ_r. These turn out to be insufficient to accurately capture the heterogeneity in returns estimated in the data. In particular, they imply that the return on wealth is too high for middle-class households, which in turn implies that these households own "too much" wealth. This is reflected in table 1.11 which compares the distribution of wealth under this specification of returns with that of the baseline model and the data. Taken together, the previous observations suggest that an economy where ex-post *luck* is the only source of return heterogeneity is unlikely to accurately describe the underlying (true) return

[36]Alternatively, the model fit can be assessed from table 1.19 in appendix 1.8.4.

Table 1.11: Wealth shares: model (one type) and data

Percentile	Model ($J = 1$)	Baseline	Data
Bottom 50%	1.5	1.5	1.5
Middle 40%	52.9	22.8	22.1
Top 10%	45.6	75.7	76.4
Top 5%	29.1	68.9	64.9
Top 1%	8.82	55.5	37.2

process faced by households. Motivated by this result and the discussion in section 1.2.6 , I consider type dependence as a way of better capturing the return heterogeneity observed in the data. Next, I discuss the choice of the total number of types.

As discussed in section 1.4.3, the cross-sectional evidence on returns to wealth is not sufficient to disentangle between ex-ante return heterogeneity and ex-post return differences, creating a challenge for the identification of the parameters associated with the return process. To check how the results depend on the relative importance of ex-ante return differences, captured by return "types", relative to ex-post return differences, I estimate the model for different numbers of types. In particular, I consider the case of two, three (baseline) and four return types.

Figure 1.11: Model fit: average return by wealth percentile

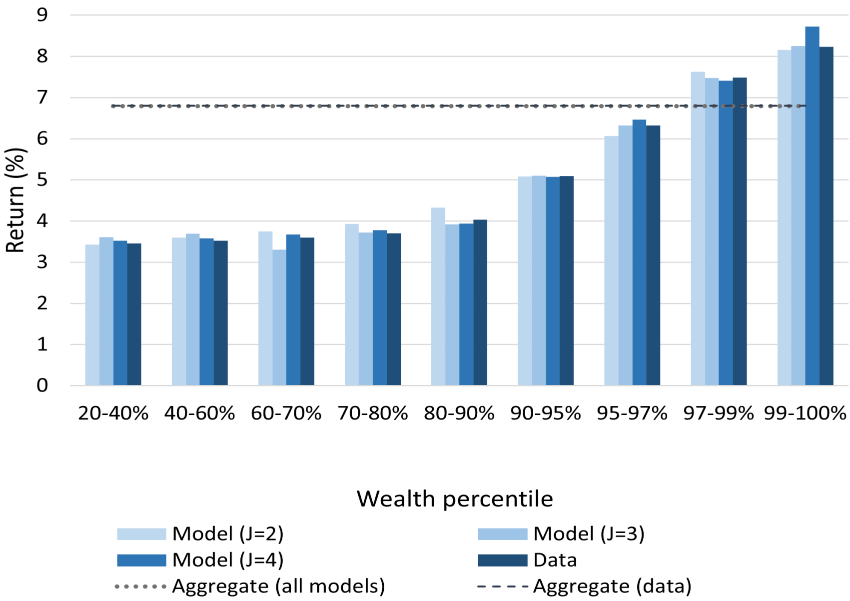

Relative to the one-type specification, considering two, three and four return types considerably improves the matching of the targeted (average) return moments. This is observable in figure 1.11 which compares the targeted return moments implied by the different models

Table 1.12: Idiosyncratic volatility of returns: alternative return specifications

Wealth percentile	Two types	Three types	Four types
20%	20.4%	6.9%	7.4%
90%	21.0%	14.5%	8.3%
99%	23.9%	5.8%	9.7%

and the data. Notice, however, that the idiosyncratic volatility parameter is different in each of these specifications. This can be seen in table 1.12. In general, the idiosyncratic volatility parameter falls with the number of individual return types. Intuitively, considering more return types allows for relatively more "ex-ante" return heterogeneity, which implies that the targeted return moments can be matched with less "ex-post" idiosyncratic return heterogeneity.

Table 1.13: Wealth shares: model and data

Percentile	Model ($J = 2$)	Model ($J = 3$)	Model ($J = 4$)	Data
Bottom 50%	1.5	1.5	1.5	1.5
Middle 40%	23.2	22.8	26.8	22.1
Top 10%	75.2	75.7	71.7	76.4
Top 5%	67.7	68.9	62.7	64.9
Top 1%	57.1	55.5	49.6	37.2

Turning to the wealth distributions implied by the different model specifications, table 1.13 summarizes the implied wealth shares of different percentiles.[37] Broadly, the distribution of wealth under the different specifications is fairly similar up to the 99th wealth percentile. This indicates that, while the individual return parameters are not precisely estimated, the implications of the estimated return heterogeneity for most of the wealth distribution do not depend substantially on the precise division of returns into ex-ante versus ex-post differences (conditional on matching the targeted moments). However, for the top 1% (and smaller groups within the top 1% such as the top 0.1% or the top 0.01%) the model-implied wealth shares vary considerably depending on the calibration. This is depicted in table 1.14. The lack of detailed panel data is an important limitation to understanding the importance of return differences for households at the very top, and the cross-sectional evidence from the SCF seems insufficient to provide a precise answer. With these caveats in mind, the different model simulations nonetheless suggest that return heterogeneity can generate large wealth shares at the very top of the distribution. This motivates further work, including data collection efforts, to better measure

[37] Additionally, appendix 1.8.4 displays the Lorenz curves associated with the different model specifications.

the wealth returns of small, but very wealthy groups of individuals.

Table 1.14: The right tail of the wealth distribution: model and data

Percentile	Model ($J = 2$)	Model ($J = 3$)	Model ($J = 4$)	Data
Top 1%	57.1	55.5	49.6	37.2
Top 0.1%	42.8	22.7	32.2	14.1
Top 0.01%	18.5	7.8	13.8	5.1

1.6 Conclusion

This paper studies the role of return heterogeneity as a driver of wealth inequality in the United States. The first contribution is to investigate the empirical relationship between returns and wealth using data on U.S. household finances. I find substantial differences in the average return on wealth, arising both from differences in portfolio allocations and return heterogeneity on similar assets. Richer households earn, on average higher returns: a gap of 4.7 percentage points between the 20th and the 99th percentile.

To understand the implications of return differences for wealth inequality, I consider return heterogeneity in a partial equilibrium model of household saving behavior and calibrate it in order to be consistent with the estimated heterogeneity in U.S. data. I find that return heterogeneity is able to explain the large wealth concentration at the top observed in the US. For example, adding return heterogeneity to the standard model of earnings risk raises the top 10% wealth share from 36% to 76%, fully matching its empirical counterpart.

This paper takes a first step towards quantifying the importance of return heterogeneity for wealth inequality in the United States, and there are many avenues which may be fruitful for future research. First, it will be important to further investigate the deeper determinants of return differences (e.g. portfolio choice, investment skill, information) and their relative importance. Relatedly, it will be interesting, as well as challenging, to quantify the distributional implications of heterogeneous returns in a general equilibrium model in which prices are determined endogenously. Finally, the data limitations that the literature still faces when it comes to measuring returns to wealth in micro data suggest that there are potentially large gains from collecting administrative data on both income and wealth. This would greatly improve our understanding of the fundamental drivers of wealth inequality worldwide.

1.7 Appendix 1: Data

1.7.1 Survey of Consumer Finances' sample design

The Survey of Consumer Finances has a complex sample design intended to accurately measure aggregate wealth in the United States. Broadly, households are selected from a double sampling procedure where, first, a sample is selected from a standard multi-stage area-probability design intended to provide good coverage of assets that are broadly distributed (for example, primary homes). Then, a second sample is selected from statistical records — the Individual Tax File — derived from tax data by the Statistics of Income Division of the Internal Revenue Service (IRS). The list provided by the IRS consists of very high income families who are selected into the second sample with the aim of disproportionately picking families that are likely to be relatively wealthy. Weights are used to combine information from the two samples to make estimates for the full U.S. population.

1.7.2 The distribution of wealth in the United States, 1989-2019

Table 1.15 depicts wealth shares in the United States between 1989 and 2019. It also includes the Gini index of wealth, which is an alternative measure of dispersion or inequality along the entire distribution.

Table 1.15: Evolution of wealth inequality in the United States, 1989-2019

	1989	1992	1995	1998	2001	2004	2007	2010	2013	2016	2019
Bottom 50%	3.0	3.3	3.6	3.1	2.8	2.6	2.5	1.2	1.1	1.2	1.5
Middle 40%	29.5	29.7	28.5	28.4	27.7	28.0	26.1	24.4	23.9	21.8	22.1
Top 10%	67.5	67.0	67.9	68.6	69.5	69.4	71.4	74.4	75.0	77.0	76.4
Top 1%	30.2	30.1	34.7	34.0	32.2	33.3	33.6	34.2	35.5	38.7	37.2
Gini index	0.792	0.787	0.791	0.801	0.805	0.809	0.816	0.846	0.850	0.860	0.852

1.7.3 Calculating the yield component of returns

This section describes, in detail, the asset classes and income flows from the SCF that were included in the calculation of returns. In each survey period I only consider households whose head is, at least, twenty years old.

Interest-earning assets. The category "interest-earning assets" includes all liquid assets (money market accounts, checking accounts, savings accounts, call accounts, prepaid cards), certificates of deposit, directly and indirectly held bonds and the cash value of life insurance. The income flow associated with these assets is the total annual interest income reported by households.

Public equity. Public equity is the sum of households' direct holdings of stock plus other public equity owned indirectly through mutual funds or other vehicles. The income flow generated by public equity is the payment of dividends to stockholders.

Private business equity. Private businesses are self-reported by households as the share of net equity in the non-publicly traded businesses owned. These include unincorporated businesses (proprietorships and partnerships) and incorporated businesses (for example, subchapter S and C-corporations). To be clear, the SCF asks households the following questions: (1) "Now I would like to ask you about businesses you may own. Do you (and your family living here) own or share ownership in any privately-held businesses, including farms, professional practices, limited partnerships, private equity, or any other business investments that are not publicly traded? Do not include corporations with publicly-traded stock or any partnerships that have already been recorded earlier."; (2) "Is it a limited partnership, another type of partnership, an LLC, a subchapter S corporation, another type of corporation, or something else?"; (3) "What could you sell your (family's) share for? What is it worth? About how much would it cost to buy a similar asset?".

To estimate the profits generated by private businesses I follow Moskowitz and Vissing-Jørgensen (2002). More specifically, the reported net income associated with each business is adjusted for corporate taxes, retained earnings and the unreported labor income of entrepreneurs.

The tax adjustment assumes a tax rate of 30%[38] on profits for C-corporations and 0% for S-corporations and partnerships. In order to get an estimate of profits that excludes earnings retained by firms, a fraction of after tax profits is deducted — 40% for C-corporations and 20% for S-corporations and partnerships. These percentages correspond to estimates of the ratio of retained earnings to after tax profits in NIPA data (National Accounts). I use the values estimated by Moskowitz and Vissing-Jørgensen (2002) and Kartashova (2014). Finally, I deduct from profits an estimate of the entrepreneur's labor income when salary is not reported. That is, I impute a salary to individuals who are self-employed, have ownership in a private company with active management interest, but report no salary. For these individuals, the reported annual hours worked are multiplied by an estimate of the wage rate of similar individuals in the survey who work in paid employment.[39] The final estimate of profits from a private business is, thus, equal to the reported earnings from the business minus the estimated corporate taxes paid, the estimated retained earnings and the labor income associated with the entrepreneur.

Real estate. The SCF collects information on the market value of all real estate owned by households (primary residence, other residential real estate and non-residential real estate). This is the value of the stock of real estate owned by families. The total net income generated

[38]This is a measure of the average effective corporate tax rate applied in the United States, following Moskowitz and Vissing-Jørgensen (2002).

[39]The wage rate is imputed based on the individual's age, educational attainment and gender.

by real estate is calculated as follows.

The SCF collects and groups together the total income that comes from "rents, royalties and trusts". From this income I deduct the amount that does not come from rents. The latter is obtained by assuming that if (1) the household does not own primary residence or any other real estate or (2) the household does not own any other real estate and has declared royalties, then the income reported as "rents, royalties and trusts" is associated with royalties or trust income but not rents. Then I assume that the remaining income after this adjustment corresponds to rental income from real estate other than the primary residence. I calculate the ratio of rental income to the gross value of other real estate, which fluctuates between 3% and 9% over the sample period. Then, I impute rents to primary residences by assuming an identical rent-to-value ratio to the one of other real estate.

Debt. The SCF subdivides households' debt into six different categories: loans secured by primary residence, debt secured by other residential real estate, other lines of credit, credit card debt, installment loans (e.g. vehicles, education) and other debt. The following assumptions were used in the calculations of interest earnings and payments.

Debt secured by primary residence. Households report the amount owed in mortgages and other lines of credit secured by their primary residence. They are asked about the current annual interest rate paid on these up to three loans. I calculate interest payments as the product of mortgage debt (on the main home) multiplied by the geometric average of the interest rates reported on the mortgage loans (if the reported amount of debt is positive).[40]

Debt secured by other residential property. Households are asked about the amount owed in mortgages or other loans related to secondary real estate and the corresponding current interest rate up to two loans. Again, I calculate interest payments by multiplying total household debt (related to secondary property) multiplied by the geometric average of the reported interest rate paid on mortgages.

Other lines of credit. Households are asked about the amount owed in other lines of credit up to three loans. As before, I calculate interest payments by multiplying the reported household debt by the geometric average of the reported interest rate paid on other lines of credit.

Credit card balances. Households are asked about the amount owed in bank accounts associated with credit cards. I calculate interest payments by multiplying household debt (related to credit card accounts) multiplied by the geometric average of the reported interest rate. Note

[40]I only use the geometric average of the interest rates reported if the household has more than one loans. If the household has only one line of credit, I use the reported interest rate paid on this loan.

that until 1992 no question asked was on the interest rate paid. Here, I assume — as in the SCF — a monthly rate of 2.5% on credit card balances.

Installment loans: vehicles, education and others. Households provide information on the amount owed related to vehicle, education or other loans and the associated interest rate paid. Interest payments are estimated as the product of the remaining debt owed and the reported interest rate, with the following considerations. In 1989 the survey asks about the amount owed in consumer loans and the corresponding interest rates paid but it does not divide them into vehicles and education loans. In 1992, I use the geometric average of reported interest rates by type of debt — vehicles, education and other (residual). For "other" installment debt, I assume the interest rate paid is a geometric average of the interest rates reported on loans for vehicles and education.

Other debt. The SCF asks households about any other debt they may have and the annual interest rate paid on those loans. Interest rate payments are calculated as the product of the interest rate and the stock of debt still owed.

1.7.4 Total return by wealth category

Table 1.16 shows the estimated aggregate returns by wealth class using the methodology described in section 1.2.3. Figure 1.12 plots the return of selected asset classes.

Table 1.16: Estimated annualized returns by wealth category, 1990-2019

	90-92	93-95	96-98	99-01	02-04	05-07	08-10	11-13	14-16	17-19	**Average**
Interest-earning assets	3.77	3.54	2.65	2.51	1.97	1.64	1.68	1.19	1.08	0.97	2.10
Public equity	5.86	9.73	25.80	1.72	-2.98	7.03	-8.78	11.80	8.70	7.79	6.67
Private businesses	9.92	15.51	21.67	9.93	14.74	15.72	-1.70	17.04	15.40	15.36	13.36
Real estate	-1.11	1.94	4.50	7.91	11.20	7.40	-3.74	5.70	10.65	8.51	5.30
Debt	4.25	3.26	3.23	2.83	2.83	2.57	2.48	1.86	1.81	1.85	2.70
Other financial	-12.96	3.55	-8.93	14.71	-0.25	9.75	-2.34	-7.94	5.73	2.60	0.39
Other nonfinancial	-4.78	9.59	1.52	5.88	4.22	-2.35	0.71	0.17	2.36	1.37	1.87

Figure 1.12: Aggregate returns by asset class

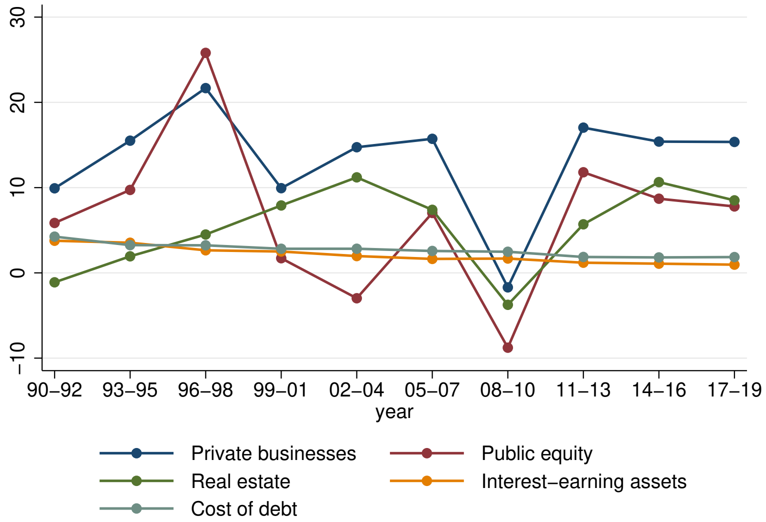

1.8 Appendix 2: Model

1.8.1 Numerical Solution

This section describes the algorithm used to numerically compute the stationary equilibrium. It is an adaptation of the methods developed in Achdou et al. (ming).

The economy can be represented by the following system of equations:

$$\begin{aligned}\rho v_j(a,y,r) =& \max_c u(c) + \partial_a v_j(a,y,r)(y+ra-c) + \partial_y v_j(a,y,r)\mu_y(y) \\ & + \partial_r v_j(a,y,r)\mu_{r,j}(r) + \frac{1}{2}\partial_{yy} v_j(a,y,r)\sigma_y^2(y) + \frac{1}{2}\partial_{rr} v_j(a,y,r)\sigma_{r,j}^2(r)\end{aligned} \tag{1.15}$$

$$\begin{aligned}0 =& -\partial_a[s_j(a,y,r)g_j(a,y,r)] - \partial_y[\mu_y(y)g_j(a,y,r)] - \partial_r[\mu_{r,j}(r)g_j(a,y,r)] \\ & + \frac{1}{2}\partial_{yy}[\sigma_y^2(y)g_j(a,y,r)] + \frac{1}{2}\partial_{rr}[\sigma_{r,j}^2(r)g_j(a,y,r)]\end{aligned} \tag{1.16}$$

$$1 = \sum_{j=1}^{J} \int_{\underline{r}_j}^{\bar{r}_j} \int_{\underline{y}}^{\bar{y}} \int_{\underline{a}}^{\infty} g_j(a,y,r)\,dadydr \tag{1.17}$$

where $j = 1,2,3$ denotes agents' return *type*. Moreover, $s_j(a,y,r) \equiv y+ra-c_j(a,y,r)$ is the saving policy function and optimal consumption satisfies $c_j(a,y,r) = (u')^{-1}(\partial_a v_j(a,y,r))$. The state constraint $a \geq \underline{a}$ gives rise to the boundary condition

$$\partial_a v_j(\underline{a},y,r) \geq u'(y+r\underline{a}) \qquad , \quad j=1,2,3 \tag{1.18}$$

To solve the model numerically, one has to define boundaries for the labor income and return processes. Let $\underline{y}$ and $\bar{y}$ denote the boundaries of the y-process. Similarly the process for r_j gets reflected at $\underline{r}_j$ and $\bar{r}_j$.[41] This gives rise to the following boundary conditions:

$$0 = \partial_y v_j(a, \underline{y}, r) = \partial_y v_j(a, \bar{y}, r) \qquad , \ j = 1, 2, 3 \tag{1.19}$$

$$0 = \partial_y v_j(a, y, \underline{r}_j) = \partial_y v(a, y, \bar{r}_j) \qquad , \ j = 1, 2, 3 \tag{1.20}$$

HJB Equation. To solve the HJB equation (1.15) I use an implicit upwind finite difference method. The HJB equation is solved separately for each return type j following the same steps which I now describe. For simplicity I omit the subscript j.

The state space (a, y, r) is discretized as follows: a_k, $k = 1, ..., K$, y_l, $l = 1, ..., L$ and r_m, $m = 1, ..., M$. I use a non-equispaced grid with 1000 points for wealth a ($K = 1000$) and equispaced grids for y ($L = 10$) and r ($M = 10$). Let $v_{k,l,m}$ denote $v(a, y, r)$, $\Delta a_k^+ = a_{k+1} - a_k$ and $\Delta a_k^- = a_k - a_{k-1}$ and so on. The derivative of v in the a dimension is approximated using an upwind scheme, i.e. using either a forward or backward difference approximation depending on the sign of the drift

$$\begin{aligned} \partial_a^B v_{k,l,m} &= \frac{v_{k,l,m} - v_{k-1,l,m}}{\Delta a_k^-} \\ \partial_a^F v_{k,l,m} &= \frac{v_{k+1,l,m} - v_{k,l,m}}{\Delta a_k^+} \end{aligned} \tag{1.21}$$

Similarly, I also use an upwind method in the y and r directions. For the second-order derivatives, I use a central difference approximation:

$$\begin{aligned} \partial_y^B v_{k,l,m} &= \frac{v_{k,l,m} - v_{k,l-1,m}}{\Delta y_l} \\ \partial_y^F v_{k,l,m} &= \frac{v_{k,l+1,m} - v_{k,l,m}}{\Delta y_l} \\ \partial_{yy} v_{k,l,m} &= \frac{v_{k,l+1,m} - 2v_{k,l,m} + v_{k,l-1,m}}{(\Delta y_l)^2} \end{aligned} \tag{1.22}$$

$$\begin{aligned} \partial_r^B v_{k,l,m} &= \frac{v_{k,l,m} - v_{k,l,m-1}}{\Delta r_m} \\ \partial_r^F v_{k,l,m} &= \frac{v_{k,l,m+1} - v_{k,l,m}}{\Delta r_m} \\ \partial_{rr} v_{k,l,m} &= \frac{v_{k,l,m+1} - 2v_{k,l,m} + v_{k,l,m-1}}{(\Delta r_m)^2} \end{aligned} \tag{1.23}$$

The discretized version of (1.15) is given by

[41] For each type j, the boundaries $\underline{r}_j$ and $\bar{r}_j$ are defined as $\mu_{r,j} \pm 1.65\sigma_{r,j}$.

$$\frac{v^{n+1}_{k,l,m} - v^n_{k,l,m}}{\Delta} + \rho v^{n+1}_{k,l,m} = u(c^n_{k,l,m}) + \partial_a v^{n+1}_{k,l,m}[y_l + r_m a_k - c^n_{k,l,m}] \\ + \mu_l \partial_y v^{n+1}_{k,l,m} + \frac{\sigma_l^2}{2}\partial_{yy} v^{n+1}_{k,l,m} + \mu_r \partial_r v^{n+1}_{k,l,m} + \frac{\sigma_r^2}{2}\partial_{rr} v^{n+1}_{k,l,m} \tag{1.24}$$

and optimal consumption is implicitely defined by

$$u'(c^n_{k,l,m}) = \partial_a v^n_{k,l,m} \tag{1.25}$$

Given an initial guess $v^n_{k,l,m}$, equation (1.24) implicitely defines $v^{n+1}_{k,l,m}$. The *upwind scheme* is the method that defines when to use a backward or a forward approximation of partial derivatives. Let $x^+ = max\{x, 0\}$ and $x^- = min\{x, 0\}$ for any scalar x. Then, the upwind finite difference approximation of the HJB equation (1.24) is given by

$$\frac{v^{n+1}_{k,l,m} - v^n_{k,l,m}}{\Delta} + \rho v^{n+1}_{k,l,m} = u(c^n_{k,l,m}) \\ + \partial_a^F v^{n+1}_{k,l,m}[y_l + r_m a_k - c^n_{k,l,m}]^+ + \partial_a^B v^{n+1}_{k,l,m}[y_l + r_m a_k - c^n_{k,l,m}]^- \\ + \mu_l^+ \partial_y^F v^{n+1}_{k,l,m} + \mu_l^- \partial_y^B v^{n+1}_{k,l,m} + \frac{\sigma_l^2}{2}\partial_{yy} v^{n+1}_{k,l,m} \\ + \mu_r^+ \partial_r^F v^{n+1}_{k,l,m} + \mu_r^- \partial_r^B v^{n+1}_{k,l,m} + \frac{\sigma_r^2}{2}\partial_{rr} v^{n+1}_{k,l,m} \tag{1.26}$$

To update $v^{n+1}_{k,l,m}$ given $v^n_{k,l,m}$ requires solving a system of linear equations. The system implied by (1.26) can be written in matrix notation as

$$\frac{1}{\Delta}(v^{n+1} - v^n) + \rho v^{n+1} = u^n + (A^n + \Lambda + \Omega) v^{n+1} \tag{1.27}$$

where Λ and Ω are the $(K \times L \times M)$ matrices that summarize the stochastic processes for income y and returns r respectively. v^{n+1} and v^n are vectors of length $(K \times L \times M)$. The system represented by (1.26) is solved iteratively using Matlab's sparse matrix routines.

Summary of the Algorithm. Guess $v^0_{k,l,m}$, $k = 1, ..., K$, $l = 1, ..., L$ and $m = 1, ..., M$. Then, for $n = 0, 1, 2, ...$ follow

1. Compute $\partial_a v^n_{k,l,m}$ using (1.21).
2. Compute c^n from (1.25), i.e. $c^n_{k,l,m} = (u')^{-1}(\partial_a v^n_{k,l,m})$.
3. Find $v^{n+1}_{k,l,m}$ from (1.26), which makes use of (1.22) and (1.23).
4. If $v^{n+1}_{k,l,m}$ is close enough to $v^n_{k,l,m}$ stop. Otherwise go back to step 1.

KF Equation. The solution to (1.16), which also have to satisfy (1.17), can be obtained by solving

$$\tilde{A}^T g_j = 0 \tag{1.28}$$

where $\tilde{A} = A^n + \Lambda + \Omega$. The matrix A^n is the one obtained from the final HJB iteration described above. Intuitively, the matrix $\tilde{A}$ summarizes the evolution of the stochastic process (a_t, y_t, r_t). To find the stationary distribution over the state, one solves the eigenvalue problem $\tilde{A}^T g_j = 0$. To solve this problem and simultaneously impose (1.17), I do as follows. First, fix $g_{k,l,m,j} = 0.1$ for $(k, l, m) = (1, 1, 1)$ (any other point works as well). Then, solve $\tilde{A}^T \hat{g}_j = 0$ and renormalize $g_j = \delta_j \hat{g}_{k,l,m,j} / (\sum_{m=1}^{M} \sum_{l=1}^{L} \sum_{k=1}^{K} \hat{g}_j \Delta a \Delta y \Delta r)$ for each $j = 1, 2, 3$. Recall that δ_j is the mass of individuals of type j and $\Delta a \equiv 0.5(\Delta a^+ + \Delta a^-)$.

1.8.2 Optimization routine

The calibration procedure is adapted from Guvenen (2016). The objective is to find the set of parameters Θ that solve

$$\hat{\Theta} = arg \min_{\Theta} (\hat{M} - M(\Theta))' W (\hat{M} - M(\Theta)) \tag{1.29}$$

where $M(\Theta)$ denotes the vector of empirical moments targeted in the calibration and $\hat{M}$ the corresponding vector of moments generated by the model.

First, I create 10000 parameter combinations which are uniform Sobol (quasi-random) and, for each of these parameter combinations, I solve the model and compute the residual given by $(\hat{M} - M(\hat{\Theta}))' W (\hat{M} - M(\hat{\Theta}))$. Then, I select a subset of these points (typically 10, ranked by objective value) and use the Nelder-Mead's downhill simplex algorithm to find the local minimum around these points. The parameter combination associated with the lowest residual (out of the ten) is a candidate solution of the global minimization problem. I then check slight variations of this parameter combination to make sure that there is no other point with lower residuals.[42] If there is such a point in the neighborhood of the initial candidate solution, then I use that point as a new guess before applying the Nelder-Mead's downhill simplex algorithm again. The parameter combination which yields the lowest residual after this procedure is the solution to the global minimization problem.

[42] I take this extra step because, in some circumstances, I have found that the previous two steps did not always find the minimum residual point (although the distance to the "new" minimum was not very large). This is likely due to the highly nonlinear nature of the optimization problem.

1.8.3 Cross-sectional standard deviation of returns: model and evidence from Nordic countries

Section 1.4.3 describes the relatioship between returns and risk implied by the idiosyncratic volatility of returns. An alternative measure of this relationship is the cross-sectional standard deviation of returns. Fagereng et al. (2020) find evidence that the standard deviation of returns tends to increase with wealth for Norwegian households. For example, the cross-sectional standard deviation of financial returns rises from roughly 6% for the 20th percentile to 12% for the 90th percentile of wealth, reaching about 17% at the right end of the wealth distribution.[43] For Sweden, Bach et al. (2020) find a non-monotonic U-shaped relationship between wealth and the standard deviation of returns. For percentiles 20%-30%, the standard deviation of returns is 24.6%. It then varies between 20% and 24% for all wealth brackets up until the 99% percentile. It then increases substantially, reaching 38% for percentile 0.01%. In the present model, the cross-sectional standard deviation of returns rises from 3.8% for the 20% percentile to 8.3% for the 90%-99% wealth bracket and, then the standard deviation falls again for the top 1%, when it is equal to 2.8%.

1.8.4 Alternative model specifications

No borrowing. To avoid targeting the bottom 50% share of the wealth distribution, I now assume that agents cannot borrow, that is, $\underline{a} = 0$. The remaining parameters and targeted empirical moments are identical to the baseline. Table 1.17 summarizes the model fit and table 1.18 shows the wealth shares implied by this model.

Table 1.17: Model fit $\underline{a} = 0$

	Model $J = 3$	Data
Aggregate return	6.8	6.8
Average returns by percentile		
20-40%	3.4	3.6
40-60%	3.5	3.7
60-70%	3.6	3.3
70-80%	3.7	3.7
80-90%	4.0	3.9
90-95%	5.1	5.1
95-97%	6.3	6.3
97-99%	7.5	7.5
99-100%	8.3	8.3

[43] Some caution is required in directly comparing the estimates mentioned in this paper. The evidence presented by Fagereng et al. (2020) regarding the relationship between returns and wealth corresponds to "financial wealth," which excludes real estate and private equity assets.

Table 1.18: Wealth shares: model and data (2019)

Percentile	Model	Data
Bottom 50%	6.8	1.5
Middle 40%	19.4	22.1
Top 10%	73.8	76.4
Top 5%	68.6	64.9
Top 1%	58.9	37.2

No type dependence. Table 1.19 summarizes the model fit of the one-return-type specification. The matching is particularly poor for the targeted aggregate rate of return (5.5% vs. 6.8%) and the average return of percentiles 60%-70%, 70%-80%, 80%-90% and 90%-95%, which are too high in the model compared to the data.

Table 1.19: Model fit one type

	Model $J = 1$	Data
Aggregate return	5.5	6.8
Wealth bottom 50%	1.5	1.5
Average returns by percentile		
20-40%	3.4	3.6
40-60%	3.6	3.7
60-70%	3.9	3.3
70-80%	4.3	3.7
80-90%	4.9	3.9
90-95%	5.8	5.1
95-97%	6.5	6.3
97-99%	7.2	7.5
99-100%	8.1	8.3

Two, three and four return types. Figure 1.13 plots the Lorenz curves associated with the different model specifications with two, three (baseline) and four return types. As discussed in the main text, the implied distribution of wealth is broadly similar for the different specifications, up to the 99% percentile of the wealth distribution.

Figure 1.13: Lorenz curve

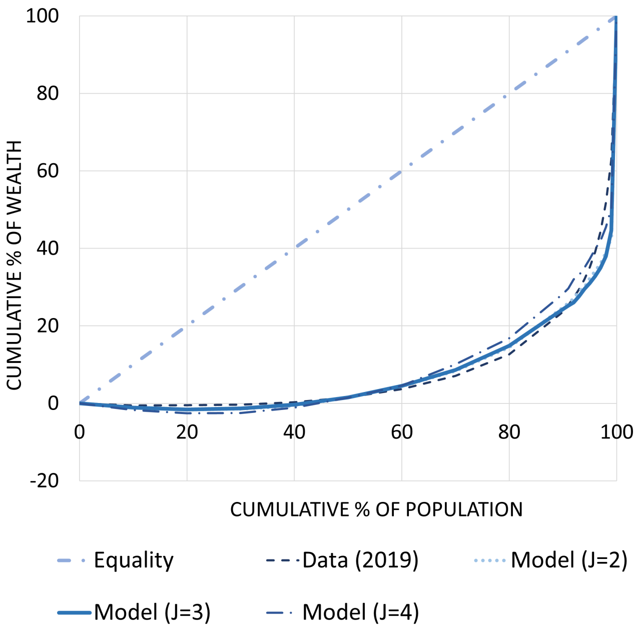

Chapter 2

Bubbles and Stagnation

Figure 2.1: Short term interest rates, 1990-2017

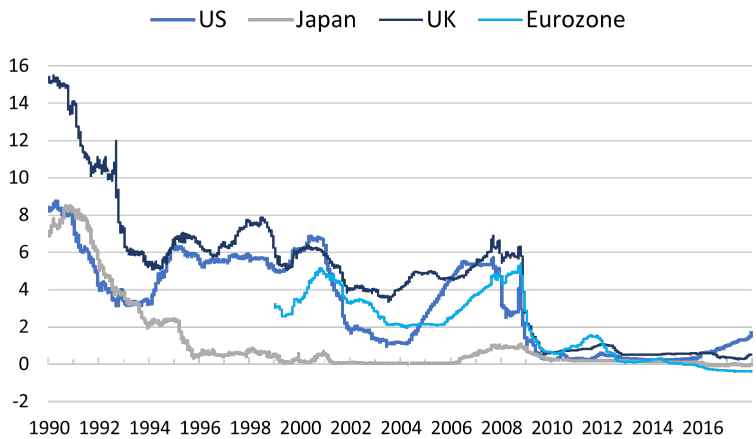

Source: ICE Benchmark Administration Limited (IBA).

2.1 Introduction

The aftermath of the Great Recession has been characterized by sluggish growth and interest rates very close to the zero lower bound (figure 2.1). The remarkably slow recoveries from the financial crisis generated fears of a new *Lost Decade*[1] and motivated the revival of the secular stagnation hypothesis. The idea behind this hypothesis is that structural factors may create a chronic excess of savings relative to the demand for new investments, depressing interest rates, output and growth. More than a decade after the crisis, the secular stagnation debate remains alive as there are still plausible concerns about chronically low growth in the industrialized world. The financial crisis may have precipitated a noticeable drop in interest and growth rates

[1]The Lost Decade became a popular term to describe the decade of economic stagnation in Japan that followed the Japanese asset price bubble's collapse in late 1991 and early 1992.

Figure 2.2: Wealth, inflation and the output gap

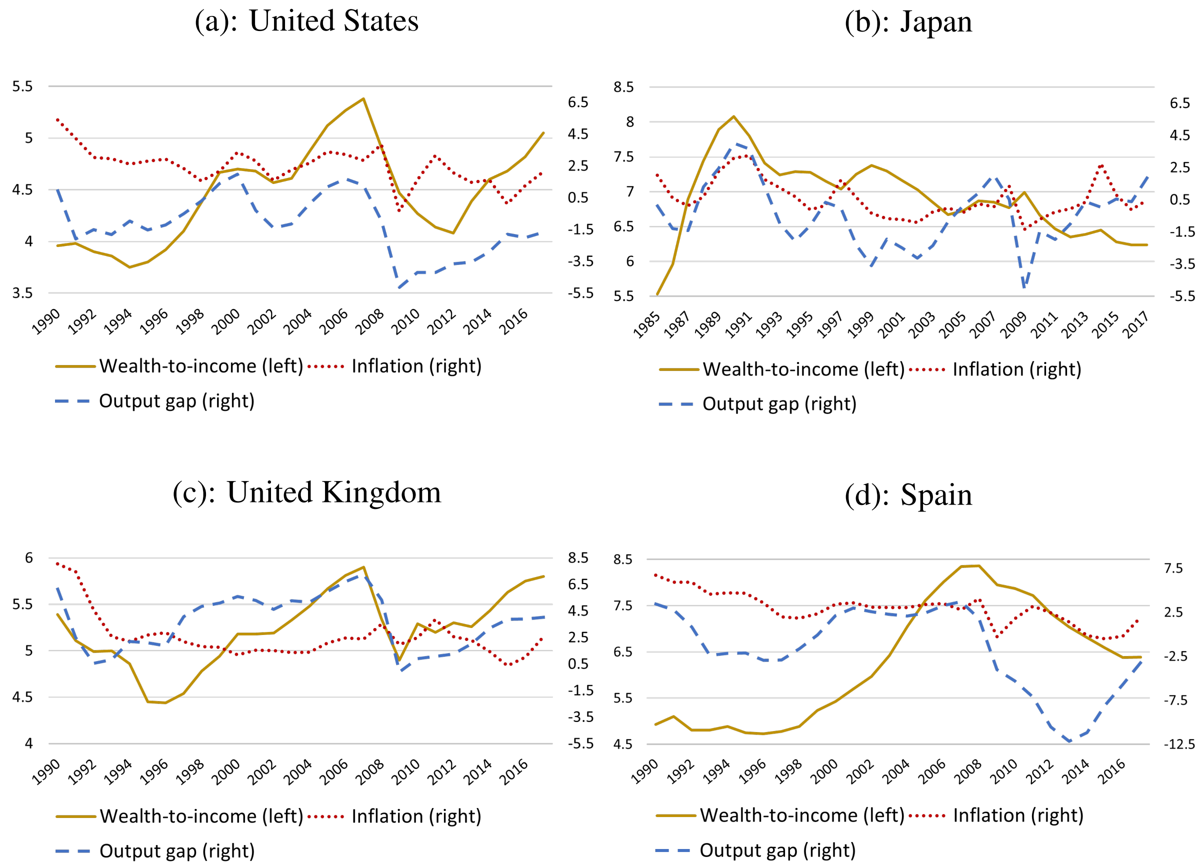

Source: The World Inequality Database (WID), FRED – Federal Reserve Bank of St. Louis, OECD.

but there is a broad consensus that natural real rates have been falling since the 1970s and that this trend is unlikely to reverse in the foreseeable future (Rachel and Summers, 2019).

At the same time, the last thirty years have been marked by recurrent episodes of large fluctuations in asset prices, often associated with "bubbles" due to an apparent disconnection between asset prices and fundamentals. Figure 2.2 is suggestive of such phenomena: it plots the aggregate wealth-to-income ratio in four advanced economies over the past three decades (solid line). In the United States, two bubbly episodes stand out as large appreciations followed by sudden collapses in net worth, reflecting the observed movements in stock and real estate prices. Similarly, the Japanese economy observed a rapid rise in stock and real estate prices in the late 1980s, which are commonly interpreted as bubbles that eventually collapsed in the early 1990s. The lower panels in figure 2.2 provide two additional examples, in the United Kingdom and in Spain, respectively. Here, asset prices appreciated rapidly during the early 2000s, but an important share of these price gains was quickly undone around 2008, coinciding with the bursting of real estate bubbles.

Motivated by the events of the past few decades, economists have speculated about the implications of asset bubbles for economies in stagnation, raising important questions about whether bubbles can affect economic outcomes and how. Lawrence Summers and Paul Krugman have provided some of the most notable commentaries offering a view that is nicely summarized by Krugman:[2] "...*how can you reconcile repeated bubbles with an economy showing no sign of inflationary pressures? Summersâs answer is that we may be an economy that needs bubbles just to achieve something near full employment â that in the absence of bubbles the economy has a negative natural rate of interest*". Despite the importance of the public debate on bubbles and the secular stagnation, this discussion has mostly been done in informal outlets without the support of a formal framework to clearly identify the mechanisms at work. The goal of this paper is to provide a model to formally address the several questions raised and inform the discussion on bubbles and the secular stagnation.

Three broad insights emerge from the analysis. First, bubbles have different implications in good and in bad times. Outside of the stagnation environment, bubbles crowd out private lending and reduce welfare by tightening borrowing constraints and distorting consumption decisions. On the contrary, bubbles can be expansionary and raise welfare in times of stagnation, when aggregate demand is chronically low and the economy operates below full capacity. In this case, bubbles absorb resources that would otherwise be wasted and crowd in aggregate consumption. Second, bubbles make stagnation less likely. On the one hand, the stagnation regime becomes less likely as bubbles raise the natural interest rate and can prevent output from falling below potential. On the other hand, bubbles provide a mechanism to get out of stagnation by expanding aggregate demand and raising equilibrium employment. Finally, bubbles that may randomly collapse are less likely to avoid stagnation than safe bubbles. However, some bubble risk may be desirable provided it is not too high. Riskier bubbles are associated with lower interest rates, which loosen borrowing constraints and may improve the allocation of resources.

To understand why bubbles can be expansionary in this environment, note that stagnation is the by-product of a severe shortage of assets that creates an excess supply of savings relative to demand. Away from the zero lower bound (ZLB), the gap between the demand and the supply of savings can be eliminated by an appropriate reduction in the nominal (and real) interest rate that raises the demand for funds (adjustment via asset supply). At the ZLB, however, the nominal interest rate cannot be lowered further and equilibrium is, instead, restored through a fall in output that reduces desired savings (adjustment via asset demand). Stagnation thus emerges from the combination of nominal rigidities, the ZLB and the structural factors that depress the natural real rate. By raising the supply of assets and the natural interest rate, bubbles lessen the need for output to fall and, therefore, make the stagnation equilibrium less likely.

The framework developed in this paper contains the key ingredients of the secular stagnation literature, on the one hand, and the literature on rational bubbles on the other. I briefly describe them next. The economy is inhabited by overlapping generations of young, middle-

[2]Quote from Krugman (2013b). See also Summers (2013) and Krugman (2013a) for other contributions to the public debate on bubbles and the secular stagnation.

aged and old. Individuals can borrow and lend to each other but financial frictions impose an upper limit on the demand for funds. This affects young agents who issue debt to middle-aged households in order to consume at early stages of life. The latter, in turn, save a part of their income to consume at old age. When financing conditions are tight, desired savings exceed the economy's ability to borrow (i.e., the young cannot issue enough assets) which triggers a fall in the natural interest rate. In the absence of bubbles, the existing scarcity of assets generates a shortfall of aggregate demand that cannot be compensated by a fall in prices due to nominal rigidities and the ZLB. In this case, equilibrium is restored with a fall in labor demand and output. However, if a bubble is created, it can prevent output from falling below potential. Essentially, bubbles allow old households to compensate for the low consumption of the young, who face tight borrowing constraints. By selling bubbles to the middle-aged, the old can absorb the excess savings generated by the tight borrowing conditions, making it no longer necessary for output to fall for an equilibrium to be reached. The model predicts that, relative to the fundamental stagnation regime, bubbles raise equilibrium output and inflation, which is consistent with the empirical behavior of inflation (dotted line) and the output gap (dashed line) during the bubbly episodes depicted in figure 2.2.

After presenting the main intuitions in the context of safe bubbles, I introduce bubble risk and discuss some of the implications of stochastic bubble dynamics. I show that bubble equilibria become less likely as the risk of collapse rises because riskier bubbles require a higher level of inflation, which may not be consistent with the inflation target. Potentially, if the risk of collapse is too high, no bubble equilibria exist. In this case, the economy settles at a fundamental equilibrium which can be the stagnation trap. While safe bubbles are more likely to avoid stagnation than risky bubbles, some risk may still be desirable provided it is not too high. Riskier bubbles are associated with lower interest rates, which facilitate borrowing for the constrained households and may improve the allocation of resources.

Related Literature. This paper relates to three broad strands of the literature. First, it contributes to the extensive work on liquidity traps. Eggertsson and Woodford (2003), Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2017) are just a few examples that feature the possibility of zero lower bound episodes. Interestingly, none of these models allows the zero lower bound to be binding in steady state, as it can only do so temporarily after a shock. In contrast, Eggertsson et al. (2019) argue that liquidity traps may last indefinitely, opening the door to a secular stagnation. In their model a permanent liquidity trap is possible as a result of a deleveraging shock, a fall in population growth, an increase in income inequality or a drop in the relative price of investment. Although my model shares many similarities with theirs, I use this framework to analyze a natural question that is ignored in their paper: how can asset bubbles affect the equilibrium of an economy that can fall permanently into a stagnation trap? Two closely related papers are Benhabib et al. (2001) and Benigno and Fornaro (2017) who also feature the possibility of permanent liquidity traps. Although the stagnation equilibrium of these models resembles, in several ways, the permanent stagnation featured in my model, there are two important differences. First, the core mechanisms behind the permanent stagnation equilibrium are different (expectations-driven vs. fundamentals-driven) and, more importantly, the two papers exclude from the analysis the possibility of asset bubbles.

Second, this work relates to the literature on asset shortages and their macroeconomic implications. A closely related work is that of Caballero and Farhi (2017), where a strong scarcity of safe assets coupled with a binding ZLB generates a permanent fall in output and employment ("safety trap"). One implication of their analysis is that while safe assets and bubbles improve a "liquidity trap", risky bubbles are dominated by safe assets for improving a safety trap. In this paper, I argue that while safe bubbles are more likely to avoid stagnation, some risk may be desirable because it relaxes financing constraints.

Finally, this paper relates to the literature on rational asset bubbles that goes back to Samuelson (1958) and Tirole (1985). It closely relates to the theoretical models of bubbles due to borrowing constraints, where this type of friction allows an intrinsically worthless asset to trade at a positive price. Some examples include Farhi and Tirole (2012), Martin and Ventura (2012, 2016) and Caballero and Krishnamurthy (2006), in which bubbles affect investment and production decisions. In my model, instead, bubbles act as aggregate demand shifters as they interact with the frictions that generate the secular stagnation: tight borrowing constraints, a ZLB on the nominal interest rate and nominal wage rigidities. Closely related is the work of Biswas et al. (2020) who also study the aggregate implications of bubbles in a setting that features a ZLB and nominal wage rigidities. However, similar to the previously mentioned examples, the paper develops an investment-driven demand for bubbles and their macroeconomic effects operate primarily through their impact on investment and production.

The rest of the paper is organized as follows. Section 2.2 outlines a model of secular stagnation with bubbles. Section 2.3 characterizes the equilibria of the model when bubbles are deterministic. Section 2.4 considers the possibility of stochastic bubble destruction and investigates how this type of risk affects the steady-state equilibria before risk is realized. Section 2.5 concludes.

2.2 A Model of Secular Stagnation and Bubbles

2.2.1 Setup

Consider an economy populated by overlapping generations of young, middle-aged and old households that live for three periods. Time is discrete and infinite, $t \in \{0, ..., \infty\}$. The middle-aged supply labor and own firms, earning income from wages and profits, denoted by W_t and Z_t respectively. I assume that labor supply from the middle-aged is inelastic and equal to $\bar{L}$. To consume at early stages of life, the young borrow from middle-aged households who, in turn, save for retirement when they use all of their income to consume. There are three assets that agents can use to shift consumption across time: a one period risk-free bond, A_t^i, $i = y, m, o$, with return r_t[3]; one period nominal debt whose return, i_t, is controlled by the government[4] ; and a bubble, B_t, which grows at rate r_t^b. Finally, I assume that there is a limit, denoted by D_t, to the amount that each generation can repay and, thus, borrow.

The bubble asset has no fundamental value but rational agents may still decide to purchase this asset if they expect to be able to resell it in the future. Moreover, bubbles may appear and collapse randomly. When the economy is in the fundamental state, a bubble may be created in the following period with probability p, i.e. $P(B_{t+1} > 0|B_t = 0) = p$, where $p \in [0, 1]$. If, instead, there is a positive bubble at time t, its value may collapse in the following period with probability λ, i.e. $P(B_{t+1} = 0|B_t > 0) = \lambda$ where $\lambda \in [0, 1]$. Individuals born in period t seek to maximize expected utility

$$\mathbb{E}_t\Big\{log(C_t^y) + \beta log(C_{t+1}^m) + \beta^2 log(C_{t+2}^o)\Big\}$$

where C_t^y is the consumption of the household when young, C_{t+1}^m its consumption when middle-aged and C_{t+2}^o consumption when old.

Perfectly competitive firms maximize period-by-period profits, taking prices as given. Let P_t denote the price level.

$$Z_t = \max_{L_t} P_t Y_t - W_t L_t \tag{2.1}$$

s.t.

$$Y_t = L_t^\alpha \tag{2.2}$$

Optimization by the firm yields the following labor demand condition

$$\frac{W_t}{P_t} = \alpha L_t^{\alpha-1} \tag{2.3}$$

Given the structure described above, each household born at time t faces the following budget constraints

$$C_t^y = -A_t^y \tag{2.4}$$

$$C_{t+1}^m = \frac{W_{t+1}}{P_{t+1}} L_{t+1} + \frac{Z_{t+1}}{P_{t+1}} + (1 + r_t)A_t^y - A_{t+1}^m - B_{t+1} \tag{2.5}$$

$$C_{t+2}^o = A_{t+1}^m(1 + r_{t+1}) + B_{t+2} \tag{2.6}$$

$$B_{t+2} = 1_{\{B_{t+1}>0\}} B_{t+1}(1 + r_{t+1}^b) + 1_{\{B_{t+1}=0\}} B_{t+2}^N \tag{2.7}$$

$$-(1 + r_t)A_t^i \leq D_t \tag{2.8}$$

[3]Note that A_t^i can be positive or negative depending on whether the agent is buying the asset (lending) or issuing it (borrowing).

[4]I assume that it trades in zero net supply. Microfoundations for the explicit introduction of money include cash-in-advance constraints or money in the utility function.

where equation (2.4) corresponds to the budget constraint of the young agent who consumes the total amount borrowed. Equation (2.5) reflects the budget constraint of the middle-aged agent who receives income from labor and profits. Moreover, she repays the amount borrowed when young, lends to the current young and buys the bubble. At old age, the individual simply consumes the (gross) return to her savings, as stated in equation (2.6). The amount of old-age consumption derived from the bubble depends on whether it already existed in the previous period or not. If the bubble was positive at $t+1$, then $B_{t+2} = B_{t+1}(1 + r^b_{t+1})$. If, instead, there was no bubble at $t+1$, a bubble of size B^N may be created at $t+2$ with probability p (equation 2.7). Finally, inequality (2.8) reflects the existence of a financial friction that imposes an upper bound on the amount that can be repaid by each household.

In what follows, I assume that inequality (2.8) binds for the young, so that they borrow as much as possible. I am, thus, considering an economy where young agents would like to borrow more in order to consume, but financing frictions impose a limit to the amount of credit they can obtain. Therefore

$$C^y_t = -A^y_t = \frac{D_t}{1 + r_t} \tag{2.9}$$

The old at time t consume all of the disposable income

$$C^o_t = A^m_{t-1}(1 + r_{t-1}) + B_t \tag{2.10}$$

The middle-aged must decide how much to consume at t and how much to save for old age. Moreover, they must decide how to allocate savings between the existing vehicles. Solving this problem requires the following optimality conditions.

$$\frac{1}{C^m_t} = \beta \mathbb{E}_t \frac{1 + r_t}{C^o_{t+1}} \tag{2.11}$$

$$\frac{1}{C^m_t} = \beta \mathbb{E}_t \frac{1}{C^o_{t+1}}(1 + i_t)\frac{P_t}{P_{t+1}} \tag{2.12}$$

$$\frac{1}{C^m_t} = \beta \mathbb{E}_t \frac{1 + r^b_t}{C^o_{t+1}} \tag{2.13}$$

Equations (2.11) and (2.12) imply the usual Fisherian relationship

$$1 + r_t = (1 + i_t)\mathbb{E}_t \frac{P_t}{P_{t+1}} \tag{2.14}$$

Moreover, equations (2.11) and (2.13) imply that if $A^m_t > 0$ and $B_t > 0$, agents must be indifferent between purchasing the two assets in the sense that expected marginal utility from C^o_{t+1} is equalized. If the growth rate of the bubble were lower than the one implied by (2.11) and (2.13), holding a bubble would not be sufficiently attractive and agents would only buy bonds. On the other hand, if it was larger, the middle-aged — the ones who supply savings in this model — would like to borrow to purchase the bubble and this cannot be an equilibrium either. In addition, the bubble must be feasible which requires enough funds to purchase it. For simplicity, I assume no growth and therefore a positive growth rate of the bubble is ruled out (i.e. a bubble is only sustainable when $r^b \leq 0$). In general, though, it is only required that the bubble does not grow faster than the economy's resources which allows for a positive bubble growth rate if income is also growing.

In the spirit of Eggertsson et al. (2019), I introduce downward nominal wage rigidities. Suppose that in any given period t the household would never accept a wage W_t lower than $\gamma W_{t-1} + (1-\gamma) P_t \alpha \bar{L}^{\alpha-1}$. Note that $\gamma = 0$ corresponds to the case with perfectly flexible wages and $\gamma = 1$ corresponds to a perfectly downward rigid wage. In this last case, households at t would not accept a wage lower than what it was at $t-1$. The degree of price rigidity is thus governed by γ.

Finally, consider a central bank that sets the nominal interest rate according to the following Taylor rule

$$1 + i_t = \max\left\{1, (1+i^*)\left(\frac{\mathbb{E}_t \Pi_{t+1}}{\Pi^*}\right)^{\phi_\pi}\right\} \quad , \ \phi_\pi > 1 \tag{2.15}$$

where $\mathbb{E}_t$ is the expectations operator conditional on information available at time t, $\Pi_{t+1} = \frac{P_{t+1}}{P_t}$, Π^* is the (exogenous) inflation target and i^* is the nominal interest rate target that is consistent with the inflation target. This rule implies that the central bank tries to stabilize expected inflation around its target Π^* unless it is constrained by the zero lower bound where $i_t = 0$. The greater is ϕ_π, the more sensitive is the central bank to deviations of inflation from target, and the less relevant is i^* for the determination of equilibrium inflation. It turns out that considering the limit $\phi_\pi \to \infty$ is useful to analyze the implications of asset price bubbles because, as discussed next, bubbles affect the natural interest rate and, thus i^*. Considering $\phi_\pi \to \infty$ allows me to study the implications of bubbles while abstracting from how the central bank updates i^* in response to the creation or the collapse of a bubble.

2.3 Deterministic bubbles

Consider first the case of deterministic bubbles. I later relax this assumption but for now bubbles do not appear randomly ($p = 0$) nor they collapse ($\lambda = 0$). This section characterizes the equilibria of the model in two different settings: (1) a perfectly flexible economy and (2) an economy with nominal rigidities.

2.3.1 Equilibrium

An equilibrium is defined as a collection of $\left\{C_t^y, C_t^m, C_t^o, A_t^y, A_t^m, B_t, L_t, Y_t, Z_t\right\}$ and prices $\left\{P_t, W_t, r_t, r_t^b, i_t\right\}$ that solve $(2.1) - (2.3), (2.5), (2.7)$ and $(2.9) - (2.15)$ given an exogenous process for $\left\{D_t\right\}$ and an initial bubble B_0. I start by deriving aggregate supply (AS) and

Figure 2.3: Aggregate supply

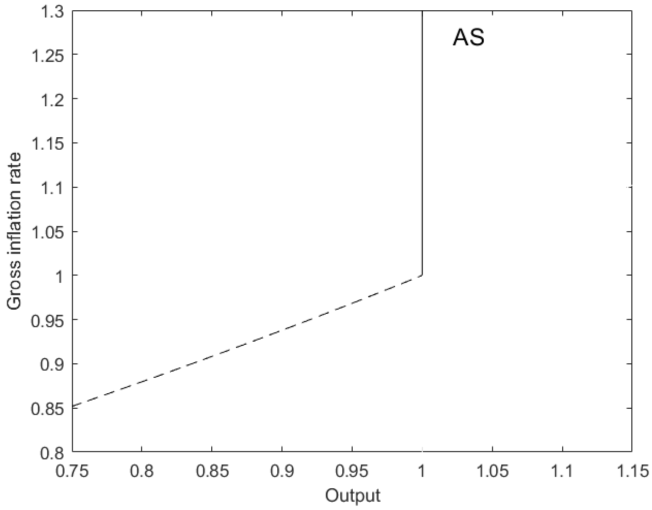

aggregate demand (AD) which must be equalized in equilibrium. For now, I ignore transitional dynamics and focus on steady state equilibria.

Aggregate supply consists of two regimes: one in which the real wage clears the market at full employment (if $\Pi \geq 1$) and another in which the lower bound on nominal wages is binding (if $\Pi < 1$). When $\Pi \geq 1$, the bound on the nominal wage does not bind and the real wage is consistent with full employment $\bar{L}$. Output is given by

$$Y = \bar{L}^{\alpha} \equiv Y^{f} \qquad \text{for } \Pi \geq 1 \tag{2.16}$$

In this case, the supply curve is represented by a vertical line determined by the supply of labor $\bar{L}$, as depicted in figure (2.3).[5]

When $\Pi < 1$, the wage norm binds and the real wage is given by $\omega = \frac{(1-\gamma)\alpha\bar{L}^{\alpha-1}}{1-\gamma\Pi^{-1}}$. In this case, aggregate supply (AS) is given by

$$Y = \left(\frac{1-\gamma}{1-\gamma\Pi^{-1}}\right)^{\frac{\alpha}{\alpha-1}} Y^{f} \qquad \text{for } \Pi < 1 \tag{2.17}$$

Production increases with inflation because of wage stickiness. As inflation rises, real wages decrease and firms demand more labor. This is represented by the upward sloping, dashed segment in figure (2.3).

Aggregate demand is given by the sum of all generations' consumption. Let $Y_t^m \equiv \frac{W_{t+1}}{P_{t+1}} L_{t+1} + \frac{Z_{t+1}}{P_{t+1}}$ and combine (2.5), (2.6) and (2.11) to get the demand for bonds from the middle-aged

$$A_t^m = \frac{\beta}{1+\beta}(Y_t^m - D_{t-1}) - B_t \tag{2.18}$$

[5]In this example, $\bar{L} = 1$.

The middle-aged households save a constant fraction of their disposable income and use it to purchase bonds and bubbles. The remaining fraction is used to consume.

$$C_t^m = \frac{1}{1+\beta}(Y_t^m - D_{t-1}) \tag{2.19}$$

Equilibrium in the bond market requires that the supply of this asset from the young equalizes its demand from the middle-aged, i.e., $\sum_i A_t^i = 0$. Given (2.18), this implies that the young's consumption is equal to

$$C_t^y = \frac{\beta}{1+\beta}(Y_t^m - D_{t-1}) - B_t \tag{2.20}$$

Finally, using the fact that $r_t = r_t^b$, old age consumption is given by

$$C_t^o = \frac{\beta}{1+\beta}(Y_{t-1}^m - D_{t-2})(1 + r_{t-1}) \tag{2.21}$$

Old households simply consume the (gross) return to their savings. Steady state aggregate demand is, thus, the sum of equations (2.19), (2.20) and (2.21)[6]:

$$Y = D + \frac{1+\beta}{\beta}\left(\frac{D}{1+r} + B\right) \tag{2.22}$$

Combining the previous expression with the Taylor rule (2.15) yields aggregate demand as a function of inflation:

$$Y = \begin{cases} D + \frac{1+\beta}{\beta}\left(D \cdot \Sigma^* \frac{1}{\Pi^{\phi_\pi - 1}} + B\right) & \text{if } i > 0 \\ D + \frac{1+\beta}{\beta}\left(D \cdot \Pi + B\right) & \text{if } i = 0 \end{cases} \tag{2.23}$$

where $\Sigma^* \equiv (1+i^*)^{-1}(\Pi^*)^{\phi_\pi}$. For a given bubble size B, equation (2.23) says that, at positive nominal rates, aggregate demand is decreasing in inflation. The intuition is straightforward: as inflation increases, the central bank raises de nominal rate more than proportionally, which raises the real rate and reduces aggregate demand. Moreover, the greater is ϕ_π, the flatter is the downward sloping segment of the AD curve. When $i = 0$, demand is increasing in inflation. At the zero lower bound, the central bank cannot affect the nominal rate as inflation rises. This leads to a fall in the real interest rate which stimulates demand. Figure (2.4) plots aggregate demand against inflation and, as in the remainder of the paper, I will focus on the case $\phi_\pi \to \infty$[7]. The AD curve has a kink at the level of inflation that corresponds to $i = 0$.

[6] An alternative way to derive aggregate demand would be to impose bond market clearing. To see this, note that (2.22) is exactly equivalent to $\frac{D}{(1+r)} = \frac{\beta}{1+\beta}(Y - D) - B$. The left-hand side corresponds to the demand for savings by the young which must be equal to the supply of savings by the middle-aged on right-hand side.

[7] For the graphical representations of aggregate demand in this paper, I set ϕ_π equal to 5000 as an approximation to the case $\phi_\pi \to \infty$.

Figure 2.4: Aggregate demand

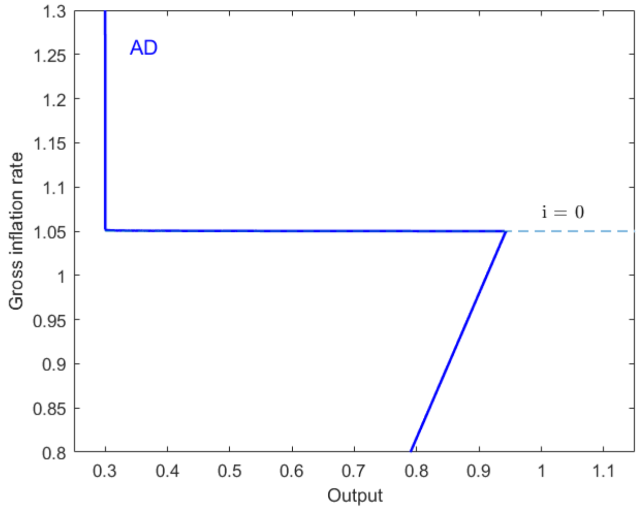

2.3.2 Flexible wage economy: $\gamma = 0$

When there are no price rigidities, the nominal wage can fully adjust so that the labor market always clears at full employment. Production is fixed at $Y^f = \bar{L}^{\alpha}$ which can be thought of as an endowment.

Fundamental equilibrium

The economy here is identical to the endowment economy of Eggertsson et al. (2019), where $B = 0$. Nonetheless, I will characterize its dynamics as it will be useful for later comparison. Imposing $B = 0$ and plugging Y^f into equation (2.22), the equilibrium of this economy can be fully characterized by

$$\frac{D}{(1+r)} = \frac{\beta}{1+\beta}(Y^f - D) \tag{2.24}$$

Graphically, this equilibrium is illustrated in figure 2.5. The left-hand side of (2.24) – the demand for savings by the young generation – corresponds to the downward sloping solid line. The vertical dashed lines depict the supply of savings from the middle-aged. The equilibrium interest rate will in general depend on the debt limit D, as well as the discount factor. Point A in figure 2.5 depicts the equilibrium in the bond market for one possible parametrization. In this example, the interest rate that equalizes the demand and the supply of savings is negative. Intuitively, this is an economy where bond supply (borrowing from the young) is strongly limited by the exogenous constraint D. At positive rates, middle-aged savings exceed the amount that the young agents can borrow. The interest rate needs to fall below zero so that the demand and the supply of savings are equalized.

Figure 2.5: Equilibrium in the bond market

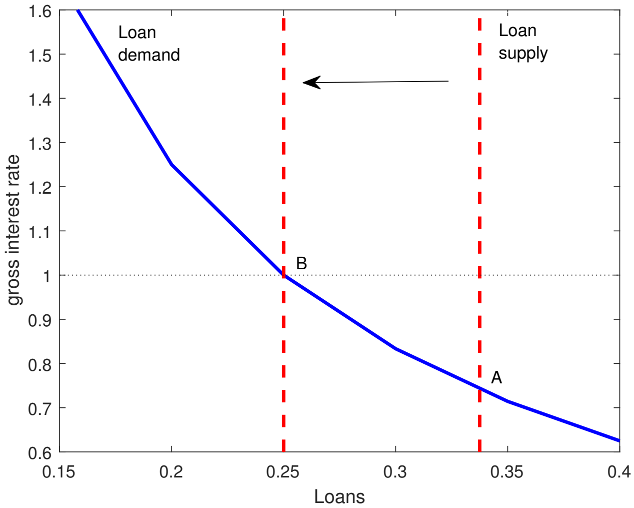

Bubble equilibria

I now consider the possibility of asset bubbles as described in the general setup, and analyze the conditions under which bubbles are possible and their effects on the equilibrium interest rate. Again, plugging Y^f into equation (2.22), we obtain the following equilibrium relationship.

$$(1+r) = \frac{D}{\frac{\beta}{1+\beta}(Y^f - D) - B} \tag{2.25}$$

Comparing (2.24) and (2.25), it is clear that a positive bubble leads to an increase in the equilibrium rate. Compared to the bubbleless economy, the supply of loans is reduced by the fact that now the middle-aged allocate some of their savings to the bubble asset. This crowding-out effect is intuitive as bubbles compete with bonds as a saving vehicle. Graphically, the supply of loans shifts to the left as depicted in figure 2.5. But when can a positive bubble be sustained in equilibrium? To answer this question, we need to analyze bubble dynamics. Given that the bubble must grow at the interest rate, its law of motion can be expressed as

$$B_{t+1} = \frac{D_t}{\frac{\beta}{1+\beta}(Y_t^m - D_{t-1}) - B_t} B_t \tag{2.26}$$

In steady state, this economy features two equilibria — one in which there is no bubble and another with a positive bubble. To see this, impose $B_{t+1} = B_t \equiv B^*$ in (2.26) which yields two solutions: $B^* = 0$ and $B^* = \frac{\beta}{1+\beta}(Y - D) - D$[8]. In the bubbleless equilibrium, $B^* = 0$ and the equilibrium interest rate is pinned down by (2.24). Instead, in the bubbly steady state (i.e. with a positive bubble), $r^* = 0$ and the size of the bubble is just enough to equalize supply

[8]Note that a positive bubble is only possible if $\frac{\beta}{1+\beta}(Y - D) > D$, which implies that the savings of the middle-aged exceed the amount of debt that the young can repay. In this case, $r < 0$ in the bubbleless equilibrium.

and demand of liquidity[9].

In this economy, output is fixed at full employment and must be all consumed. As such, bubbles do not affect aggregate consumption but they promote a reallocation of consumption across generations. The two steady state equilibria differ in the equilibrium interest rate, bubble size and, therefore, in the amount that each generation consumes. The young consume more in the bubbleless equilibrium since a lower interest rate relaxes their financing constraint, allowing them to borrow more. On the contrary, a lower interest rate means lower returns at old age, which reduces old age consumption. The roles reverse at the bubble equilibrium, where a higher interest rate implies lower consumption for the young and larger returns for the old. Given the preferences of the middle-aged, who consume a constant fraction of their disposable income independently of the prevailing interest rate, bubbles do not affect this generation's consumption level.

What are, then, the welfare implications of asset bubbles? In this endowment economy, bubbles reduce expected lifetime utility and, thus, welfare. Essentially, output and aggregate consumption are fixed at potential, but different bubble sizes reallocate consumption between young agents and old. In particular, a larger bubble raises the real interest rate, which tightens the borrowing constraint of young individuals who are forced to consume less. Since output must be all consumed, the old agents will compensate for the fall in consumption of the young by selling the bubble to the middle-aged agents and consuming the proceedings. However, this reallocation of consumption is associated with a fall in expected lifetime utility as shown in appendix 2.6.1. Intuitively, this is an economy in which the optimal consumption profile requires a negative real interest rate due to tight financing constraints. In this environment, bubbles are sustainable and, in equilibrium, they crowd-out private lending and raise the real interest rate. This forces young agents to postpone consumption to old age in a way that is not optimal.

Nominal rigidities, the ZLB, output and bubbles

The zero lower bound on the nominal interest rate imposes a restriction on the rate of inflation. In particular, the requirement that $(1+i) \geq 1$ implies that $\Pi \geq (1+r)^{-1}$ for an equilibrium with constant inflation to exist[10]. At positive real rates, this restriction has little empirical relevance as it requires a level of inflation well below central banks' common targets. However, if the real interest rate is negative, this condition implies a positive inflation rate, possibly above the usual 2% target followed by many central banks. As the real interest rate falls, steady state inflation needs to be higher for an equilibrium to exist. But what if the central bank is unwilling to raise inflation enough? The result may be a permanent drop in output if there are some price rigidities. This is the message of Eggertsson et al. (2019) who show that if the central bank

[9] To get the equilibrium rate r^*, plug $B^* = \frac{\beta}{1+\beta}(Y - D) - D$ into equation (2.25).

[10] Note that when $i = 0$, $(1 + r) = \Pi^{-1}$.

refuses to accept enough inflation, the economy can settle at a secular stagnation equilibrium with some unemployment. And what is the role of bubbles in this setting? For a given level of output, bubbles raise the real interest rate by reducing the supply of savings in the bond market, thereby relaxing the requirement imposed on inflation by the zero lower bound. As I introduce nominal wage rigidities in the model, I will show that bubbles can prevent output from falling by generating additional liquidity and thereby raising the natural interest rate. Contrary to the flexible wage economy, bubbles can be expansionary and welfare improving in this setting. I turn to this case next.

2.3.3 Nominal wage rigidities: $\gamma \in (0, 1)$

I now assume that nominal wages are downwardly rigid, i.e. that $\gamma \in (0, 1)$. We will see that sticky wages may imply some equilibrium unemployment if the real wage cannot fall enough.

Fundamental equilibrium

To analyze the fundamental equilibrium of this economy, again impose $B = 0$. Aggregate supply is given by expressions (2.16) and (2.17), while aggregate demand is described by equation (2.23) with $B = 0$. The inflation target is assumed to be positive, i.e. $\Pi^* > 1$. Aggregate supply and demand curves are represented graphically in figure 2.6. The solid black line is the AS curve with a kink at $\Pi = 1$. The two demand curves, AD_H and AD_L, depict aggregate demand for two different sets of parameters. In particular, they differ in the choice of D which is larger for the AD_H curve.

Figure 2.6: Steady state aggregate demand and supply curves

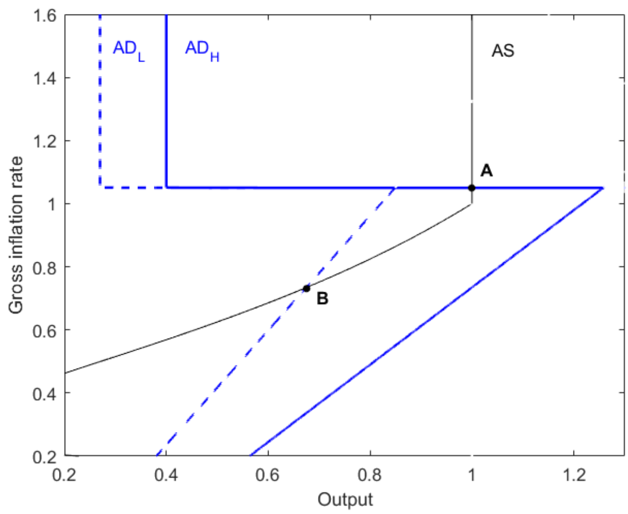

Point A depicts the equilibrium of the economy when D is *high*. In this equilibrium, the natural interest rate — the rate at which demand and supply of savings equalize at full employment — is positive. Moreover, it features positive inflation, determined by the inflation target, and zero unemployment. Now suppose that financing conditions become tighter and agents are forced to deleverage (D falls permanently)[11]. If the negative shock on aggregate demand is sufficiently large, the economy may land at a stagnation trap as depicted by point B. This is an equilibrium with permanent unemployment and a binding ZLB. As the young agents are forced to deleverage, bond supply falls. The middle-aged households, who want to save for old age, face an increased scarcity of assets that drives down the natural interest rate. Outside of the ZLB, a negative shock that reduces asset supply is compensated by a fall in the nominal (and real) interest rate which, again, relaxes financing conditions and stimulates consumption. At the ZLB, however, the interest rate cannot fall enough to stimulate demand and, instead, equilibrium in the bond market is restored endogenously by a drop in output. When the income of the middle-aged falls, asset demand also drops, which reequilibrates the market.

The stagnation equilibrium is thus a result of a severe shortage of assets coupled with a binding ZLB and nominal rigidities. When the ZLB becomes binding, the adjustment no longer happens through the interest rate but, instead, through output. As I discuss next, the key role of bubbles is to avoid the need for a drop in output by expanding the supply of assets and, in this way, avoiding stagnation.

Bubble equilibria

Consider again the possibility of trade in bubbles. As before, equilibrium output and inflation are determined by the intersection of aggregate supply and aggregate demand. However, we must also consider the circumstances under which bubbles can arise in this economy. Since the bubble must grow at the real interest rate, we can use (2.14) and (2.15) to express bubble dynamics as

$$B_{t+1} = \begin{cases} (1+i^*)\Pi_{t+1}^{\phi_\pi - 1}(\Pi^*)^{-\phi_\pi} \cdot B_t & \text{if } i > 0 \\ B_{t+1} = \Pi_{t+1}^{-1} \cdot B_t & \text{if } i = 0 \end{cases} \tag{2.27}$$

where I have assumed perfect foresight in order to eliminate the expectations operator. In order to find the steady state, impose $B_{t+1} = B_t = B^*$, which yields

$$\Pi = \begin{cases} \Pi^* & \text{if } i > 0 \\ \Pi = 1 & \text{if } i = 0 \end{cases} \tag{2.28}$$

in the limit case $\phi_\pi \to \infty$.

Consider an economy whose unique fundamental equilibrium is the stagnation trap (point A in figure 2.7). The two expressions in (2.28) correspond to additional equilibria of the model if agents are allowed to trade bubbles. The stagnation steady state is still an equilibrium of the

[11] More generally, changes in D capture movements in any fundamental force that may affect that natural interest rate such as demographics, technological progress or income inequality. See Eggertsson et al. (2019) for a detailed description of the different mechanisms.

Figure 2.7: Steady state aggregate demand and supply with bubbles

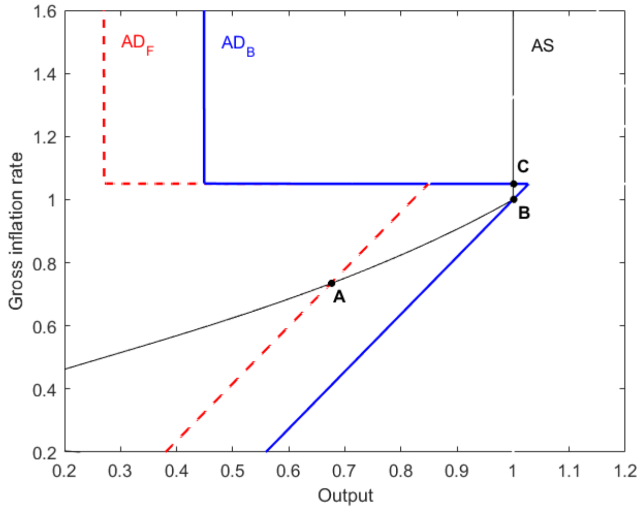

model if the agents coordinate on the bubbleless steady state ($B^* = 0$). However, now there are two additional equilibria that feature a positive bubble, represented by points B and C. Both are associated with full employment and a zero interest rate, i.e. $r^* = 0$. Moreover, the size of the bubble is identical and just enough to equalize supply and demand of liquidity. Points B and C only differ in the nominal interest and inflation rates that are consistent with the central bank's policy rule.

At the stagnation equilibrium, financing frictions impose an upper limit on the demand for funds (supply of bonds). At this point, the demand for assets by the middle-aged households exceeds the amount of assets that the economy is able to produce (debt issued by the young). The ZLB prevents the interest rate from adjusting enough such that the supply of assets (and the consumption) of the young is restored. Without any other asset in the economy, equilibrium is reestablished through a drop in output that, by reducing income, decreases the demand for bonds by the middle-aged. The key role played by bubbles in this economy is to avoid the drop in output by reequilibrating demand and supply of liquidity through an increase in asset supply. Once the old are allowed to sell bubbles to the middle-aged, the supply of assets rises, equalizing its demand at full employment. The old are, thus, compensating for the fall in the consumption of the young that follows a tightening of their financing constraint. As they provide additional liquidity, bubbles are expansionary and sustain higher levels of consumption and employment. As shown in appendix 2.6.2, all generations are unambiguously better off at equilibria B and C since all agents can consume more. In the presence of nominal rigidities, safe bubbles are welfare improving by allowing the economy to avoid the stagnation equilibrium and increasing the consumption of all generations. In this sense, bubbles can be desirable in times of stagnation.

2.4 Stochastic bubbles

Section 2.3 discussed the implications of bubbly episodes for economies in stagnation under the assumption that bubbles are safe assets. In practice, however, bubbly episodes are generally associated with large and volatile movements in asset prices that are difficult to predict. In this section, I consider the possibility of stochastic bubble destruction and discuss the implications of this type of risk for the steady-state equilibria. I focus on the equilibria *before* bubble risk is realized.

2.4.1 Equilibrium

Consider an economy that is at steady state with a positive bubble. In any future period, the existing bubble may collapse with probability $\lambda \in (0, 1]$ and I assume that, if it does, it remains equal to zero forever. Optimization by the middle-aged households implies the following optimality conditions:

$$\frac{1}{C_t^m} = \beta(1 + r_t)\left[\frac{\lambda}{A_t^m(1 + r_t)} + \frac{1 - \lambda}{A_t^m(1 + r_t) + B_t(1 + r_t^b)}\right] \tag{2.29}$$

$$\frac{1}{C_t^m} = \beta(1 + r_t^b)\left[\frac{1 - \lambda}{A_t^m(1 + r_t) + B_t(1 + r_t^b)}\right] \tag{2.30}$$

They imply that middle-aged households must be indifferent between investing in the bond or in the bubble asset, giving rise to the following expression for the steady-state risk premium

$$\frac{r^b - r}{1 + r} = \frac{\lambda}{1 - \lambda}\frac{B(1 + r^b) + D}{D} \tag{2.31}$$

That is, stochastic bubbles are associated with a risk premium that is greater when the bubble is riskier (greater λ) or when it is larger relative to D. In a steady-state equilibrium, the net return on the bubble conditional on no collapse is equal to zero, $r^b = 0$, so that (2.31) implies a decreasing relation between the riskless rate and the bubble: a larger bubble raises the risk premium which implies a lower risk-free rate.

On the other hand, the optimality condition (2.30) and bond market clearing imply an increasing relation between the interest rate r and the bubble B,

$$\left[1 + \frac{\beta(1 - \lambda)}{1 + r}\right]D + \left[1 + \beta(1 - \lambda)\right]B = \beta(1 - \lambda)(Y - D) \tag{2.32}$$

Intuitively, a higher interest rate lowers the value of other assets, which requires a larger bubble to absorb the supply of savings for any given income Y.

The equilibrium interest rate r and the equilibrium bubble B can be found by combining equations (2.31) and (2.32), which yield

$$1 + r = \frac{D}{\frac{\beta}{1+\beta}\lambda(Y - D) + D} \tag{2.33}$$

$$B = (1 - \lambda)\frac{\beta}{1 + \beta}(Y - D) - D \tag{2.34}$$

In equilibrium, the size of the steady state bubble shrinks with its riskiness: the greater the risk of collapse, the larger is the bubble risk premium and the lower is the riskless rate. This raises the value of the non-bubble assets and reduces the size of the equilibrium bubble.

2.4.2 Stochastic bubbles and secular stagnation

To understand the aggregate implications of stochastic bubbles in a secular stagnation, it is useful to analyze their impact on the natural interest rate. Similar to the deterministic case, a positive stochastic bubble raises the natural interest rate relative to the fundamental state, making the stagnation equilibrium less likely. However, the risk of collapsing affects the feasibility of bubbles and, therefore, their potential to expand economic activity. The intuition is as follows. As is clear from (2.33) and (2.34), a riskier bubble is smaller and is associated with a lower natural interest rate than a safe bubble. A lower natural interest rate raises the level of inflation that is required for a steady state full employment equilibrium to exist[12]. If the inflation target is too low relative to this required level of inflation, no bubble equilibria exist. In other words, full employment equilibria become less likely as the risk of collapse rises because they require a higher level of inflation, which eventually may not be consistent with the inflation target. Potentially, if the risk of collapse is too high, no bubble equilibria exist. In this case, the economy settles at a fundamental equilibrium which may well be the stagnation trap.

The previous discussion suggests that stochastic bubbles are less likely to avoid stagnation than safe bubbles. In this sense, safe bubbles are preferred to stochastic bubbles. Interestingly, however, some risk may be desirable provided it is not too high. To understand the intuition behind this result, consider the set of bubbles that are feasible and consistent with full employment. These are, thus, bubbles that are not "too risky" relative to the central bank's inflation target. Within this set of bubbles, those that have a higher risk of collapsing imply a lower riskless rate which, in turn, relaxes the borrowing constraints of the young households. In appendix 2.7.2, I show that this raises welfare whenever

$$\lambda < \frac{1}{1 + \beta^2}\Big[1 - \beta(1 + \beta)\frac{D}{Y^f - D}\Big]. \tag{2.35}$$

Intuitively, for low values of λ, the equilibrium bubble is large and the riskless rate is high. When borrowing constraints are tight, welfare can be raised by redistributing consumption from the old savers to the young borrowers, who are consuming too little. As λ rises, the

[12]Recall, from the discussion at the end of section 2.3.2, that the ZLB imposes a lower bound on the rate of inflation that is consistent with steady state equilibria.

equilibrium bubble shrinks and the old consume less, while the young households borrow and consume more. This raises welfare up to the point at which λ exceeds the threshold implied by (2.35). For even risker bubbles, i.e. λ larger, the previous redistribution from the old to the young is no longer welfare improving. Above the threshold, the old become the ones who consume too little compared to the optimal consumption profile, such that greater risk becomes undesirable[13].

2.5 Concluding Remarks

The model developed in this paper provides a framework to think about the role of asset bubbles in stagnant economies. Here, a severe shortage of assets depresses the interest rate and output. Low interest rates open the door for the emergence of bubbles that, in this setting, can be expansionary as they increase the supply of assets and prevent output from falling below its potential level. For this to happen, however, bubbles cannot be too risky. If the risk of collapsing is too large, bubbles are no longer consistent with full employment and they fail to expand economic activity. Although highly stylized, this model allows us to focus on the role of bubbles in the type of economies vulnerable to stagnation, and provides a formalization of the much discussed relationship between bubbles and the secular stagnation hypothesis.

Even though growth recovered in the major developed economies after the Great Recession, the pace of global economic activity remains sluggish. Moreover, the structural forces that explain the steady decline in interest rates over the past decades are still at work (e.g. the aging of these economies). Close to the ZLB the effectiveness of conventional monetary policy in counteracting negative demand shocks is limited. In this paper, I show how sufficiently safe bubbles provide a mechanism to sustain aggregate demand and a higher level of employment in such an economy. This type of mechanism, however, depends crucially on the ability to coordinate expectations, which opens the door to recessions triggered by coordination failures.

The theory developed in this paper is consistent with the narrative that, before the collapse of the housing bubble that was associated with the Great Recession, bubbles were sustaining higher levels of employment and interest rates. After the collapse of the bubble, several advanced economies entered a stagnation trap, with low interest rates and higher unemployment, and central banks constrained by the ZLB. This suggests that, while bubbles can be expansionary, they can also push the economy into a slump when they burst. The instability associated with bubbly episodes suggests that policy may be desirable to deal with the drawbacks of this instability. Further research on the interactions between asset bubbles, demand and economic activity will be useful to shed light on this question.

[13]Recall that this comparison is between bubble equilibria conditional on a given level of output. In this case, bubble risk affects the distribution of consumption across the different generations due to its impact on the riskless rate. In addition to this channel, bubble risk can also affect the level of equilibrium output, as discussed in the first part of section 2.4.

2.6 Appendix 1: Deterministic bubbles

2.6.1 Flexible wage economy: bubbles and welfare

In this section I compare the two steady states of the flexible wage economy — depicted in figure (2.5)— and analyze the welfare implications of asset bubbles. I compare lifetime utility in both steady states — X and Y — and show that, whenever a positive bubble is possible, it is welfare reducing. I focus on the case where $\frac{\beta}{1+\beta}(Y^f - D) > D$ so that bubbles are possible. Otherwise, the only possible equilibrium would be the bubbleless one.

Let us start by considering the bubbleless equilibrium — point X. Here, $B^* = 0$ and each generation's steady state consumption is given by

$$C^y = \frac{\beta}{1+\beta}(Y^f - D) \tag{2.36}$$

$$C^m = \frac{1}{1+\beta}(Y^f - D) \tag{2.37}$$

$$C^o = D \tag{2.38}$$

Given the agents' preferences, the lifetime utility associated with this consumption pattern is given by

$$U^X = log\Big[\frac{\beta}{1+\beta}(Y^f - D)\Big] + \beta \cdot log\Big[\frac{1}{1+\beta}(Y^f - D)\Big] + \beta^2 \cdot log(D) \tag{2.39}$$

Let us now turn to equilibrium Y with a positive steady state bubble given by $B^* = \frac{\beta}{1+\beta}(Y^f - D) - D$ and $r^* = 0$. In this equilibrium, each generation's consumption is given by

$$C^y = D \tag{2.40}$$

$$C^m = \frac{1}{1+\beta}(Y^f - D) \tag{2.41}$$

$$C^o = \frac{\beta}{1+\beta}(Y^f - D) \tag{2.42}$$

Note that, compared to equilibrium X, the middle-aged consume the exact same amount, but the young agent swaps her consumption level with the old individual. Expected lifetime utility is given by

$$U^Y = log(D) + \beta \cdot log\Big[\frac{1}{1+\beta}(Y^f - D)\Big] + \beta^2 \cdot log\Big[\frac{\beta}{1+\beta}(Y^f - D)\Big] \tag{2.43}$$

By comparing (2.39) and (2.43) we can investigate which equilibrium yields higher lifetime utility. Whenever $\frac{\beta}{1+\beta}(Y^f - D) > D$ [14], the bubbleless equilibrium is welfare superior to the bubbly one.

$$U^X - U^Y = (1-\beta^2) \cdot log\left[\frac{\frac{\beta}{1+\beta}(Y^f - D)}{D}\right] > 0 \tag{2.44}$$

2.6.2 Nominal wage rigidities: bubbles and welfare

In this section I compare the stagnation equilibrum — point A in figure (2.7) — to the two bubbly equilibria B and C.

Since they are identical to equilibrium Y in the previous section, let us first consider equilibria B and C. In both equilibria, output is at full employment, Y^f, and each generation's consumption s given by

$$C^y = D \tag{2.45}$$

$$C^m = \frac{1}{1+\beta}(Y^f - D) \tag{2.46}$$

$$C^o = \frac{\beta}{1+\beta}(Y^f - D) \tag{2.47}$$

In the stagnation equilibrium output falls below full employment. Recall that in this equilibrium, output is given by $Y = \left(\frac{1-\gamma}{1-\gamma\Pi^{-1}}\right)^{\frac{\alpha}{\alpha-1}} Y^f$ and $\Pi < 1$. Since, at this equilibrium, $(1+r) = \Pi^{-1}$, it must also be that $(1+r) > 1$. Each generation's consumption is given by

$$C^y = \frac{D}{1+r} < D \tag{2.48}$$

$$C^m = \frac{1}{1+\beta}(Y - D) < \frac{1}{1+\beta}(Y^f - D) \tag{2.49}$$

$$C^o = D \tag{2.50}$$

where I used the fact that $(1+r) = \frac{D}{\frac{\beta}{1+\beta}(Y-D)}$ to get (2.50). Since I am considering the case $\frac{\beta}{1+\beta}(Y^f - D) > D$, it becomes clear that all generations consume more in the bubbly equilibria than in the stagnation one. The bubbly equilibria are, thus, associated with higher lifetime utility than the stagnation equilibrium.

[14] And $\beta \in (0,1)$.

2.7 Appendix 2: Stochastic bubbles

2.7.1 Steady-state equilibria

Here I solve for the stochastic steady state before the realization of bubble destruction uncertainty. The economy starts at steady state with a positive bubble which may collapse in any future period. If it does, it remains equal to zero forever. I focus on the case in which the stagnation trap is the unique fundamental equilibrium.

If the bubble collapses at $t+1$, the equilibrium at $t+2$ is given by $Y_{t+2} \equiv Y^L < Y^f$, $\Pi_{t+2} \equiv \Pi^L < 1$. At $t+1$ aggregate demand is given by

$$Y^d_{t+1} = D + \frac{1+\beta}{\beta}\left[D \cdot \frac{\Pi^L}{1+i_{t+1}}\right] \tag{2.51}$$

which depends on i_{t+1}. If $(1+i_{t+1}) = (1+i^*)\left(\frac{\Pi^L}{\Pi^*}\right)^{\phi_\pi} > 1$, then $Y^d_{t+1} > Y^f$[15]. If, instead, $i_{t+1} = 0$, aggregate demand is given by $Y^d_{t+1} = Y^L$. In equilibrium, aggregate demand must be equal to supply and, thus, $i_{t+1} > 0$ cannot be an equilibrium since supply cannot exceed Y^f. Therefore, the equilibrium path must feature $i_{t+1} = 0$ and $Y_{t+1} = Y^L < Y^f$ which, together with the wage rule, imply that $\Pi_{t+1} < \dfrac{\omega_t}{\alpha \bar{L}^{\alpha-1}}$. Then, at time t (before uncertainty is realized), the middle-aged decide how much to consume and save — in bonds or the bubble — by solving the following problem

$$\max_{C^m_t, C^o_{t+1}} \mathbb{E}_t\{log(C^m_t) + \beta log(C^o_{t+1})\} \tag{2.52}$$

subject to

$$C^m_t = Y_t - D_{t-1} - A^m_t - B_t \tag{2.53}$$

$$C^o_{t+1} = A^m_t(1+r_t) + 1_{\{B_{t+1}>0\}} \cdot B_t(1+r^b_t) \tag{2.54}$$

Solving this problem implies the first-order conditions

$$\frac{1}{Y_t - D_{t-1} - A^m_t - B_t} = \beta(1+r_t)\left[\frac{\lambda}{A^m_t(1+r_t)} + \frac{1-\lambda}{A^m_t(1+r_t) + B_t(1+r^b_t)}\right] \tag{2.55}$$

$$\frac{1}{Y_t - D_{t-1} - A^m_t - B_t} = \beta(1+r^b_t)\left[\frac{1-\lambda}{A^m_t(1+r_t) + B_t(1+r^b_t)}\right] \tag{2.56}$$

The remaining equations that characterize the fundamental steady state are given by

[15] This follows from the assumptions that $\gamma > 0$ and $\Pi^* \geq 1$

$$(1 + r_t) = (1 + i_t)\mathbb{E}_t\Pi_{t+1}^{-1} \tag{2.57}$$

where $\mathbb{E}_t\Pi_{t+1} = \lambda \cdot \Pi_{t+1}^F + (1 - \lambda) \cdot \Pi_t$ and $\Pi_{t+1}^F = \eta \cdot \dfrac{\omega_t}{\alpha\bar{L}^{\alpha-1}}$, $\eta < 1$. The parameter η determines the level of inflation at $t + 1$ to which individuals coordinate in case the bubble collapses.

$$\frac{D_t}{1 + r_t} = A_t^m \tag{2.58}$$

$$1 + i_t = \max\left\{1, (1 + i^*)(\Pi^*)^{-\phi_\pi}\left[\lambda \cdot \Pi_{t+1}^F + (1 - \lambda)\Pi_t\right]^{\phi_\pi}\right\} \tag{2.59}$$

$$Y_t = (\alpha \cdot \omega_t^{-1})^{\frac{\alpha}{1-\alpha}} \tag{2.60}$$

$$\omega_t = \max\left\{\gamma\omega_{t-1}\Pi_t^{-1} + (1 - \gamma)\alpha\bar{L}^{\alpha-1}, \alpha\bar{L}^{\alpha-1}\right\} \tag{2.61}$$

$$B_{t+1} = (1 + r_t^b)B_t \tag{2.62}$$

To obtain the bubbly steady state, I solve the system of equations (2.55)-(2.62) for $Y_t = Y$, $\Pi_t = \Pi$, $A_t^m = A^m$, $B_t = B$, $i_t = i$, $r_t = r$, $r_t^b = r^b$ and $\omega_t = \omega$.

2.7.2 Bubbles, risk and welfare

Consider the set of bubbles that are feasible and consistent with full employment bubble equilibria. In steady state, the consumption of each generation is given by

$$C^y = \lambda\frac{\beta}{1 + \beta}(Y^f - D) + D \tag{2.63}$$

$$C^m = \frac{1}{1 + \beta}(Y^f - D) \tag{2.64}$$

$$C^o = (1 - \lambda)\frac{\beta}{1 + \beta}(Y^f - D) \tag{2.65}$$

Given households' preferences, the lifetime utility associated with this consumption profile is given by

$$U(\lambda) = log\left[\lambda\frac{\beta}{1 + \beta}(Y^f - D) + D\right] + \beta \cdot log\left[\frac{1}{1 + \beta}(Y^f - D)\right] + \beta^2 \cdot log\left[(1 - \lambda)\frac{\beta}{1 + \beta}(Y^f - D)\right] \tag{2.66}$$

Taking the derivative of (2.66) with respect to λ yields

$$\frac{\partial U(\lambda)}{\partial \lambda} = \frac{\beta}{1+\beta}(Y^f - D)\left[\frac{1}{C^y} - \frac{\beta^2}{C^o}\right] \tag{2.67}$$

This expression is positive whenever $\beta^2 C^y < C^o$, which is equivalent to the condition that $\lambda < \frac{1}{1+\beta^2}\left[1 - \beta(1+\beta)\frac{D}{Y^f - D}\right]$.

Bibliography

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2021, Forthcoming). Income and wealth distribution in macroeconomics: A continuous-time approach. *The Review of Economic Studies*.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Bach, L., Calvet, L. E., and Sodini, P. (2020). Rich pickings? Risk, return, and skill in household wealth. *American Economic Review*, 110(9):2703–2747.

Benhabib, J. and Bisin, A. (2018). Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature*, 56(4):1261–91.

Benhabib, J., Bisin, A., and Luo, M. (2019). Wealth distribution and social mobility in the US: A quantitative approach. *American Economic Review*, 109(5):1623–47.

Benhabib, J., Bisin, A., and Zhu, S. (2011). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica*, 79(1):123–157.

Benhabib, J., Bisin, A., and Zhu, S. (2015). The wealth distribution in bewley models with capital income risk. *Journal of Economic Theory*, 159:459–515.

Benhabib, J., Schmitt-Grohe, S., and Uribe, M. (2001). The perils of taylor rules. *Journal of Economic Theory*, 96(1):40–69.

Benigno, G. and Fornaro, L. (2017). Stagnation traps. *The Review of Economic Studies*, 85(3):1425–1470.

Bewley, T. F. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory*, 16(2):252–292.

Biswas, S., Hanson, A., and Phan, T. (2020). Bubbly recessions. *American Economic Journal: Macroeconomics*, 12(4):33–70.

Caballero, R. and Farhi, E. (2017). The safety trap. *The Review of Economic Studies*, 85(1):223–274.

Caballero, R. and Krishnamurthy, A. (2006). Bubbles and capital flow volatility: Causes and risk management. *Journal of monetary Economics*, 53(1):35–53.

Cagetti, M. and De Nardi, M. (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy*, 114(5):835–870.

De Nardi, M. and Fella, G. (2017). Saving and wealth inequality. *Review of Economic Dynamics*, 26:280–300.

Eggertsson, G. and Krugman, P. (2012). Debt, deleveraging, and the liquidity trap: a fisher-minsky-koo approach. *The Quarterly Journal of Economics*, 127(3):1469–1513.

Eggertsson, G. and Woodford, M. (2003). The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 2003(1):139–211.

Eggertsson, G. B., Mehrotra, N. R., and Robbins, J. A. (2019). A model of secular stagnation: Theory and quantitative evaluation. *American Economic Journal: Macroeconomics*, 11(1):1–48.

Epland, J. and Kirkeberg, M. I. (2012). Wealth distribution in Norway. *Evidence from a new register-based data source. Oslo, Norway: Statistics Norway*.

Fagereng, A., Guiso, L., Malacrino, D., and Pistaferri, L. (2020). Heterogeneity and persistence in returns to wealth. *Econometrica*, 88(1):115–170.

Farhi, E. and Tirole, J. (2012). Bubbly liquidity. *The Review of Economic Studies*, 79(2):678–706.

Gabaix, X., Lasry, J.-M., Lions, P.-L., and Moll, B. (2016). The dynamics of inequality. *Econometrica*, 84(6):2071–2111.

Guerrieri, V. and Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, 132(3):1427–1467.

Guvenen, F. (2016). Quantitative economics with heterogeneity: An a-to-z guidebook. *Princeton University Press*, 26(51):155.

Guvenen, F., Kambourov, G., Kuruscu, B., Ocampo-Diaz, S., and Chen, D. (2019). Use it or lose it: Efficiency gains from wealth taxation. Technical report, National Bureau of Economic Research.

Hubmer, J., Krusell, P., and Smith Jr, A. A. (2020). Sources of US wealth inequality: Past, present, and future. In *NBER Macroeconomics Annual 2020, Vol. 35*. University of Chicago Press.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969.

İmrohoroğlu, A. (1992). The welfare cost of inflation under imperfect insurance. *Journal of Economic Dynamics and Control*, 16(1):79–91.

Kartashova, K. (2014). Private equity premium puzzle revisited. *American Economic Review*, 104(10):3297–3334.

Krugman, P. (2013a). Bubbles, regulation, and secular stagnation. *The Conscience of a Liberal*, 25.

Krugman, P. (2013b). Secular stagnation, coalmines, bubbles, and larry summers. *New York Times*, 16:2013.

Kuhn, M., Schularick, M., and Steins, U. I. (2020). Income and wealth inequality in america, 1949–2016. *Journal of Political Economy*, 128(9):3469–3519.

Martin, A. and Ventura, J. (2012). Economic growth with bubbles. *American Economic Review*, 102(6):3033–58.

Martin, A. and Ventura, J. (2016). Managing credit bubbles. *Journal of the European Economic Association*, 14(3):753–789.

Moskowitz, T. J. and Vissing-Jørgensen, A. (2002). The returns to entrepreneurial investment: A private equity premium puzzle? *American Economic Review*, 92(4):745–778.

Nirei, M. and Aoki, S. (2016). Pareto distribution of income in neoclassical growth models. *Review of Economic Dynamics*, 20:25–42.

Quadrini, V. (2000). Entrepreneurship, saving, and social mobility. *Review of Economic Dynamics*, 3(1):1–40.

Rachel, L. and Summers, L. H. (2019). On secular stagnation in the industrialized world. Technical report, National Bureau of Economic Research.

Saez, E. and Zucman, G. (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, 131(2):519–578.

Samuelson, P. (1958). An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of Political Economy*, 6:467–482.

Shiller, R. J. (2015). *Irrational exuberance (3rd ed.)*. Princeton University Press.

Summers, L. (2013). Why stagnation might prove to be the new normal. *Financial Times*, 15:12.

Tirole, J. (1985). Asset bubbles and overlapping generations. *Econometrica*, 53(6):1499–1528.