

Matching Wealth Moments with Heterogeneous Returns

Decory Edwards*

2026-04-01

Abstract

Empirical evidence shows that individuals face widely varying returns on their assets. This heterogeneity has important implications for modeling wealth inequality. I incorporate heterogeneity in returns into a standard heterogeneous-agent model, and find the distribution required to match the empirical wealth distribution. Introducing return heterogeneity yields a much better match to the wealth moments measured using Survey of Consumer Finances (SCF) data than assuming homogeneous returns. In particular, the bottom 60% of the wealth distribution is matched well enough that the aggregate marginal propensity to consume will be aligned with modest estimates. In addition, it has been noted that the equivalence between a wealth tax and capital income tax fails when individuals no longer earn the same return. I compare a wealth tax and a capital income tax set to raise equal revenue. A capital income tax reduces wealth inequality more by targeting high-return households, but each policy entails different welfare trade-offs across the return distribution. There are a number of final exercises, such as estimating a lognormal distribution of returns across households and interpreting the returns heterogeneity required by the model as arising from banking-sector frictions (e.g., incomplete deposit rate pass-through).

*I am indebted to Christopher Carroll, Nick Papageorge, and John Huston, Kelly Lyons, Harry Wallace, Teresa Morrison, and the rest of the McNair Scholar Program and Economics Department at Trinity University for funding support conducting research. For helpful comments and suggestions, I thank my doctoral advisors Christopher Carroll, Nick Papageorge, and Stelios Fourakis. I would like to thank Matthew White, Alan Lujan, Mateo Velasquez-Giraldo, and Will Du, as well as my fellow cohort members Adrian Monninger, Kyung Woong Koh, and David Osten for comments and feedback as well. This work makes use of the HA toolkit made available by the Econ-Ark project.

⁰blank

1 Introduction

The unequal distribution of wealth is a well documented phenomenon across many countries. Disparities have not only persisted over time but intensified in recent years. In 2018, the aggregate wealth of the nation’s three richest individuals surpassed that of the entire bottom 50% of the U.S. wealth distribution. This point is stressed in a recent article based on Forbes magazine rankings (Institute for Policy Studies, November 2022)¹.

The unequal distribution of wealth has also long been a focus of study across disciplines. The statistics literature, for instance, has linked the distribution of income to the observable skewness in wealth distribution, and work in economics established microfoundations for wealth accumulation over the life cycle. Macroeconomic studies have examined the distribution of wealth among households to understand how the economy as a whole responds to aggregate fiscal shocks, as well as the decomposition of those effects by group (like wealth level).

The macroeconomics literature has undergone significant changes in recent years, with the widespread adoption of models that abandon the traditional representative agent assumption in their analysis. As this setting will require that in equilibrium all agents hold the same level of wealth, it is not a desirable laboratory in terms of producing model objects, like the distribution of wealth, that can be compared to real world counterparts.

Researchers began incorporating an exogenously determined income process that generates a distribution of income among households. The baseline form of (ex-post) heterogeneity in the model is the description of a permanent and transitory process for income. To account for business cycle dynamics, one can further assume that individuals face some level of potential unemployment in each period, creating a precautionary savings motive for consumers. Given that such uncertainty cannot be fully insured against, the availability of a riskless asset that partially insures against income risk results in households choosing to hold different levels of market resources optimally.

Krusell and Smith 1998’s seminal work suggests that models assuming heterogeneity in individual income perform well in matching the aggregate capital stock but poorly in matching the distribution of wealth. The next step is to assume there is some ex-ante heterogeneity among households, leading more households to optimally hold lower levels of wealth.² Carroll et al. 2017 adopt this approach and assume that agents differ in their time preferences, which reflects implicit characteristics of households relevant to their lifetime wealth accumulation. The authors find that this assumption of modest heterogeneity in time preferences is sufficient to match both the shape and skewness of the empirical distribution of wealth.

¹See Inequality.org articles data November 21, 2022: “Wealth Inequality in the United States” and “Updates: Billionaire Wealth, U.S. Job Losses and Pandemic Profiteers” (date accessed: March 27, 2023)

²Kaplan and Violante 2022’s recent work provides a comprehensive survey of incomplete markets models with heterogeneous agents featuring (i) uninsurable idiosyncratic income risk, (ii) a precautionary savings motive, and (iii) an endogenous wealth distribution.

The time preference factor (β) is one of the key parameters that influences an individual's equilibrium target level of market resources, but it is not directly observable. Estimating would require one to collect data through surveys or other methods that provide direct information from households. In contrast, differences in the rate of return on financial assets across households is possible, as this variable *is* observable.

In this paper, I provide further evidence of the heterogeneous agent framework's ability to match wealth moments by adding a single source of heterogeneity across households beyond the realization of ex-post shocks to their income. I allow households to differ in the return earned on assets first in the infinite horizon setting. I further extend the model to allow for rich life cycle dynamics. From there, I compare the effects of revenue-equivalent wealth and capital income tax policies on wealth inequality and welfare. Additionally, I interpret the heterogeneity in the returns on safe assets earned by households in the context of the transmission channel of monetary policy and its variation across the banking sector. As a final step, I rerun the entire exercise under the assumption that the returns across households are lognormally distributed.

2 Literature Review

2.1 Explaining Inequality in the Distribution of Wealth

Although wealth inequality gets a considerable amount of media coverage in modern times, it has been the object of deep reflection by academic and political thinkers for centuries. For this reason, the literature on wealth inequality is rich and has been studied by several disciplines.

Benhabib and Bisin 2018 review the literature on the documented skewness in the distribution of wealth with historical accounts of the origins of the shape of the wealth distribution. The authors then provide the traditional theoretical explanations of this unequal distribution: (i) skewness in the (exogenous) distribution of earnings, (ii) stochastic returns to wealth and savings, and, importantly, (iii) microfoundations for the evolution of wealth resulting from the consumption and saving behavior of households³.

De Nardi and Fella 2017 survey the literature on the microfoundations of wealth accumulation. A number of possible models of household consumption and saving behavior marked by observable differences between households (beyond demographic differences) lead to greater inequality in wealth accumulation over time. Earnings and rate of return risk, ex-ante heterogeneity in preferences, medical expenses, bequest motives, and entrepreneurship are all potential mechanisms which influence the shape of the distribution of wealth.

Gabaix et al. 2016 introduces the concept of speed of convergence to explain the observed evolution of income inequality over time, particularly within

³As explored in the next section, the emergence of heterogeneous agent models has been a significant development in investigating this issue. Bewley 1983, Aiyagari 1994, and Huggett 1993 are among the earliest examples.

the upper tail of the distribution over the past 40 years. Notably, they show that, replicating the empirical dynamics of inequality requires incorporating two additional forms of heterogeneity in the income process for households than are included in the standard consumption and saving models.⁴ The first form is *type dependence* in the income growth rate distribution, which accounts for some households having a higher average income growth rate. Second, *scale dependence* captures the fact that higher income levels are more susceptible to shocks to their income growth. The authors find that the former does a good job of explaining the rapid rise in income inequality, and the latter can generate infinitely fast transitions in inequality.

2.2 Measurements of heterogeneous rates of return

The rationale behind incorporating heterogeneity in rates of return to asset holdings lies in the use of novel datasets in recent empirical research to quantify the differences in returns among individuals. Using 12 years of Norwegian tax records, Fagereng et al. 2020 document the heterogeneity in realized returns. These differences occur both across individuals (*type dependence*), within and across asset classes with varying levels of risk, and within wealth deciles, where returns are positively correlated with wealth (*scale dependence*). Moreover, they find that heterogeneous returns exhibits significant persist over time and are positively correlated across generations. These findings support the assumption of ex-ante heterogeneous rates of return in the buffer-stock savings model of households, and provide a benchmark for comparing the distribution of rates of return to the empirical distribution of wealth, as in Carroll et al. 2017.

Bach, Calvet, and Sodini 2018 use administrative panel data on the balance sheets of Swedish residents to gauge historical and expected returns, as well as risks associated with asset holdings. Their analysis of portfolio performance supports the finding that heterogeneous returns substantially shape in the levels and growth of top wealth shares over time.

Campbell, Ramadorai, and Ranish 2019 examine equity holdings in India between 2002 and 2011 and find that heterogeneity in investment returns arising from the inherent randomness associated with risky assets and differences in investment strategies, is a key driver of rising inequality in portfolio holdings during the time period. The authors attribute the scale dependence of equity returns to the tendency of smaller accounts to be less diversified than larger ones.

Deuflhard, Georgarakos, and Inderst 2018 analyze household savings account investments within a heterogeneous agent, incomplete markets model with a precautionary saving motive and a *single asset* to partially insure against risk. They document substantial type dependence in the rate of return to these safe assets and attribute the heterogeneity in returns to differences in financial sophistication. Notably, explaining differences in investment returns for households

⁴Note that, although this analysis is about the distribution of income, this literature notably asserts that the distribution of wealth inherits some of its skewness from the distribution of income

is a vital step toward endogenizing this form of ex-ante heterogeneity among households in future models.

Altmejd, Jansson, and Karabulut 2024 provide causal evidence that financial education leads to significant differences in portfolio returns. Using university application records from the Swedish National Archives and data from the Swedish Income and Wealth registry, they show that individuals marginally admitted to business or economics programs hold more money in stocks and earn higher raw returns on these holdings than those not admitted.

2.3 Recent heterogeneous agent models with varying rates of return

Several studies extend the heterogeneous agent framework to allow for households to earn different returns on their assets. Of the models with a single, riskless asset and a partial equilibrium analysis of a distribution of returns across agents, my paper is unique in that it generates a realistic mass of agents below the 60-th percentile of the wealth distribution. This is important because, as we will see, this region of the wealth distribution is important for generating a reasonable estimate of the aggregate marginal propensity to consume. That said, I will discuss models that go beyond the modeling choices made in my paper.

Daminato and Pistaferri 2024 incorporate heterogeneous returns into the solution of a model of consumption-saving for households. They use data from the PSID to document heterogeneity in returns, which they find is comparable to the returns distribution measured using the Norwegian registry by Fagereng et al. 2020.

Benhabib, Bisin, and Luo 2019 propose an overlapping-generations model that incorporates intergenerational wealth transfers, with agents facing uncertainty regarding both labor and capital income. In an earlier work (Benhabib, Bisin, and Luo 2017) the authors more explicitly define household preferences for bequests to the next generation. Both studies conclude that the distribution of earnings and differences in rates of savings and bequests are crucial for accurately replicating the tails of the observed wealth distribution.

Guler, Kuruscu, and Robinson 2022 develop a life-cycle model with endogenous heterogeneity in the rate of return by incorporating households' optimal housing and mortgage decisions. Using this framework, the authors investigate the effects of aggregate fiscal shocks, including one-time stimulus payments and mortgage debt relief programs.

Menzio and Spinella 2025 introduce search frictions in financial markets within a standard infinite-horizon macroeconomic model. A distribution of returns across households arises endogenously in their model, and they use the empirical findings from Fagereng et al. 2020 regarding the distribution of returns to net worth as notable targeted moments.

2.4 Estimates of the aggregate MPC

A substantial empirical literature measures households’ marginal propensities to consume (MPCs), consistently finding that the aggregate MPC is far larger—and far more heterogeneous—than implied by representative-agent benchmarks. Across quasi-experimental settings, estimated MPCs typically fall in the range of 0.2 to 0.4. For example, Johnson, Parker, and Souleles 2006 show that households spent roughly 20–40 percent of the 2001 tax rebates in the quarter they were received, while Parker et al. 2013 find similar magnitudes for the 2008 stimulus payments. Structural and covariance-based approaches deliver comparable estimates: Blundell, Pistaferri, and Preston 2008 imply quarterly MPCs of about 0.2–0.3, and meta-analytic evidence from Havránek and Sokolova 2020 concludes that most studies cluster in this same range. At the same time, research using administrative or survey data highlights that MPCs can be substantially higher among households with low wealth or low liquidity, often exceeding 0.5 (see Jappelli and Pistaferri 2014a; Fagereng, Holm, and Natvik 2021). Taken together, these findings indicate that empirically plausible MPCs span a wide but well-documented interval.

The model developed in this paper produces MPCs that fall squarely within these empirically established ranges, even though the MPC distribution is not directly targeted in the calibration. This serves as an additional, independent validation of the framework, demonstrating that a model calibrated primarily to match wealth inequality can nevertheless replicate realistic consumption responses across the wealth distribution.

2.5 My contributions

This paper contributes to the literature in at least two ways. First, I model the labor income uncertainty as a random walk, rather than an AR(1) process. The AR(1) specification implies less uncertainty in household earnings over the life cycle, leading to less accumulation of wealth over the life cycle. By using a the permanent income framework, I account for *as much* of the dispersion of wealth across households as possible with labor income uncertainty. The remaining dispersion in wealth across households that cannot be explained by differences in earnings is attributed to heterogeneity in returns, ultimately leading to modest estimates of differences in returns across households.

Second, the life-cycle version of my model is much richer in its calibration of earnings and mortality rates than other studies. Specifically, I use the earnings profile of Cagetti 2003 that distinguishes mean earnings not only by age but by education cohort. I use age-education-dependent mortality rates from Brown, Liebman, and Pollet 2007. This approach allows households to be distinguished by both age and education, providing an additional mechanism for explaining the dispersion in wealth holdings. However, the impact of this mechanism is limited, and other sources of ex-ante heterogeneity among households, such as time preferences or the rate of return, are still needed to match wealth moments precisely. Said differently, life-cycle dynamics and labor income uncertainty are

not enough to generate a reasonably skewed wealth distribution.

3 Model

The heterogeneous agent macro literature provides a class of models whose output is a distribution of wealth which can be easily compared to data on the wealth holdings across individuals in terms of moments. This section will introduce the model I solve and simulate computationally in order to do this particular comparison of wealth moments.

3.1 Defining the stochastic income process

Each household's income (y_t) during a given period depends on three main factors: aggregate wage rate (W_t) that all households in the economy face, the permanent income component (p_t), which represents an agent's present discounted value of human wealth, and the transitory shock component (ξ_t) which reflects the potential risks that households may face in receiving their income payment during that period. Thus, household income can be expressed as:

$$y_t = p_t \xi_t W_t.$$

The level of permanent income for each household is subject to a stochastic process. In line with the labor income process described by Friedman 1957, I assume that this process follows a geometric random walk, which can be written as:

$$p_t = p_{t-1} \psi_t.$$

The white noise permanent shock to income with a mean of one is represented by ψ_t , which is a significant component of household income. The probability of receiving income during a given period is determined by the transitory component, which is modeled to reflect the potential risks associated with becoming unemployed. Specifically, if the probability of becoming unemployed is $\mathcal{U}$, the agent will receive unemployment insurance payments of $\mu > 0$. On the other hand, if the agent is employed, which occurs with a probability of $1 - \mathcal{U}$, the model allows for tax payments τ_t to be collected as insurance for periods of unemployment. The transitory component is then written as:

$$\xi_t = \begin{cases} \mu & \text{with probability } \mathcal{U}, \\ (1 - \tau_t) l \theta_t & \text{with probability } 1 - \mathcal{U}, \end{cases}$$

where l is the time worked per agent and the parameter θ captures the white noise component of the transitory shock.

3.2 Decision problem for households

Next, I present the baseline version of the household's optimization problem for consumption-savings decisions, assuming no ex-ante heterogeneity. In this case, each household aims to maximize its expected discounted utility of consumption $u(c) = \frac{c^{1-\rho}}{1-\rho}$ by solving the following:

$$\max \mathbb{E}_t \sum_{n=0}^{\infty} (\beta \mathcal{D})^n u(c_{t+n}).$$

Note that this setting follows a perpetual youth model of buffer stock savings, similar to the seminal work of Krusell and Smith 1998. To solve this problem, I use the Bellman equation, which means that the sequence of consumption functions $\{c_{t+n}\}_{n=0}^{\infty}$ associated with a household's optimal choice over a lifetime must satisfy⁵

$$\begin{aligned} v(m_t) &= \max_{c_t} u(c_t(m_t)) + \beta \mathcal{D} \mathbb{E}_t [\psi_{t+1}^{1-\rho} v(m_{t+1})] \\ \text{s.t.} & \\ a_t &= m_t - c_t(m_t), \\ k_{t+1} &= \frac{a_t}{\mathcal{D} \psi_{t+1}}, \\ m_{t+1} &= (1 + r_t) k_{t+1} + \xi_{t+1}, \\ a_t &\geq 0. \end{aligned}$$

4 Results

The estimation of the model proceeds in several steps. I begin by analyzing an infinite-horizon version of the model, in which households face uninsurable labor income risk and a borrowing constraint, and solve the consumption-saving problem using the endogenous grid method as implemented in **HARK**. When there is no heterogeneity in returns, I find the capital-to-output ratio necessary to match the capital-to-output target of $\frac{K}{Y} = 3$. Quantitative macroeconomic models calibrated to U.S. data commonly target an annual capital-to-output ratio in the range of roughly 2.5 to 3.0. For example, Lee, Luetticke, and Ravn 2021 calibrate their heterogeneous-agent model with financial frictions to match a capital-output ratio of 2.5, while Suen 2012 targets an annual capital-output ratio of 3.0 in a heterogeneous-agent model with preference heterogeneity.

When heterogeneity is present, there is a candidate distribution of returns and I simulate a large population of households forward to obtain the stationary wealth distribution. The parameters governing the return distribution are then chosen via Simulated Method of Moments (SMM) to minimize the distance

⁵Here, each of the relevant variables is normalized by the level of permanent income ($c_t = \frac{C_t}{p_t}$, and so on), following the standard state-space reduction of the problem for numerical tractability.

between simulated and empirical wealth shares from the Survey of Consumer Finances. Following existing literature, the 20th, 40th, 60th, and 80th percentiles of the wealth distribution from the 2004 wave of the survey are used as empirical targets.

After establishing the results in the infinite-horizon setting, I extend the model to a life-cycle framework with age- and education-specific income profiles, mortality risk, and initial conditions. The model is then re-solved and re-estimated when heterogeneity is and is not present.

4.1 The infinite horizon setting

4.1.1 The model with no returns heterogeneity

To solve and simulate the model, I follow the calibration scheme described in table 1.

Description	Parameter	Value	Source
Time discount factor	β	0.99 ⁴	Den Haan, Judd, and Juillard 2010
CRRA	ρ	1	Den Haan, Judd, and Juillard 2010
Capital share	α	0.36	Den Haan, Judd, and Juillard 2010
Depreciation rate	δ	0.025	Den Haan, Judd, and Juillard 2010
Time worked per employee	ℓ	1/.09	Den Haan, Judd, and Juillard 2010
Wage rate	W	2.37	Den Haan, Judd, and Juillard 2010
Unempl. insurance payment	μ	0.15	Den Haan, Judd, and Juillard 2010
Probability of survival	β	(1 - 0.00625) ⁴	Yields 40-year working life
Std. dev of log $\theta_{t,i}$	σ_θ^2	0.010 x 4 x $\sqrt{4}$	Carroll 1992, Carroll, Slacalek, and Tokuoka 2015
Std. dev of log $\psi_{t,i}$	σ_ψ^2	0.010 x 4/11 x $\sqrt{4}$	Carroll 1992, Debacker et al. 2013, Carroll, Slacalek, and Tokuoka 2015
Unemployment rate	$\bar{u}$	0.07	Mean in Den Haan, Judd, and Juillard 2010

Table 1: Parameter values (annual frequency) for the perpetual youth model.

The solution of the model with no heterogeneity in returns (the R-point model) is to find the value for the rate of return R which minimizes the distance between the models capital-to-output ratio $\frac{K}{Y}$ and the calibrated $\frac{K}{Y}$ ratio value of 3. The estimation procedure finds this optimal value to be $R = 1.0602$.

4.1.2 Incorporating heterogeneous returns

To estimate the model with heterogeneity in returns, I follow a procedure similar to the one outlined by Carroll et al. 2017. Specifically, I assume that different types of households have a time preference factor drawn from a uniform distribution over the interval $(\bar{R} - \nabla, \bar{R} + \nabla)$, where ∇ represents the level of dispersion. I then simulate the model to estimate $\bar{R}$ and ∇ so that the resulting wealth distribution aligns with the observed distribution in terms of inequality. Practically, I solve the following minimization problem:

$$\{\bar{R}, \nabla\} = \arg \min_{R, \nabla} \left(\sum_{i=20,40,60,80} (w_i(R, \nabla) - \omega_i)^2 \right)^{\frac{1}{2}}$$

subject to the constraint that the aggregate capital-to-output ratio in this model matches the calibrated value 3. Note that w_i and ω_i represent the porportion of total aggregate net worth held by the top i percent in the model and

in the data, respectively. The estimation procedure yields the optimal values of $R = 1.0204$ and $\nabla = 0.06833$ which pin down the estimated uniform distribution. The Lorenz curve associated with the wealth distribution from the model and the SCF data are compared in figure 1.




Figure 1: Comparison of R-Point and R-Dist Models.

The 45 degree line represents perfect equality. Lorenz curves further away from this line represent more unequal distributions of wealth. As we can see from the figure, the Lorenz curve generate by allowing for heterogeneous returns does a much better job at matching the shape of the SCF wealth distribution than when everyone earns the same return. Additionally, the following table compares the relevant moments of the empirical and simulated wealth distributions.

Table 2: Wealth Distribution (PY)

	0–20	20–40	40–60	60–80	80–95	Top 5%
Wealth share (data)	-0.002	0.011	0.044	0.118	0.255	0.574
Wealth share (no het.)	0.030	0.085	0.142	0.224	0.285	0.233
Wealth share (het.)	0.005	0.019	0.042	0.096	0.251	0.588

4.2 Incorporating life cycle dynamics into the model

Assumptions about the age and education level of households can have important implications for the income and mortality process of households. Next I extend the model to incorporate these life-cycle dynamics.

Households enter the economy at time t aged 24 years old with an education level $e \in \{D, HS, C\}$, an initial permanent income level $\mathbf{p}_0$, and a capital stock k_0 . The life cycle version of household income is given by:

$$y_t = \xi_t \mathbf{p}_t = (1 - \tau) \theta_t \mathbf{p}_t,$$

where $\mathbf{p}_t = \psi_t \bar{\psi}_{es} \mathbf{p}_{t-1}$ and $\bar{\psi}_{es}$ captures the age-education-specific average growth factor. Households that have lived for s periods have permanent shocks drawn from a lognormal distribution with a mean of 1 and a variance of $\sigma_{\psi_s}^2$ and transitory shocks drawn from a lognormal distribution with a mean of $\frac{1}{\varphi}$ and a variance of $\sigma_{\theta_s}^2$ with probability $\mathcal{X} = (1 - \mathcal{U})$ and μ with probability $\mathcal{U}$.

The normalized version of the age-education-specific consumption-saving problem for households is given by

$$\begin{aligned} v_{es}(m_t) &= \max_{c_t} u(c_t(m_t)) + \beta \mathcal{D}_{es} \mathbb{E}_t[\psi_{t+1}^{1-\rho} v_{es+1}(m_{t+1})] \\ \text{s.t.} \\ a_t &= m_t - c_t, \\ k_{t+1} &= \frac{a_t}{\psi_{t+1}}, \\ m_{t+1} &= (1 + r_t)k_{t+1} + \xi_{t+1}, \\ a_t &\geq 0. \end{aligned}$$

The additional parameters necessary to calibrate the life-cycle version of the model are given in table 3. The age-education dependent mean income levels come from Cagetti 2003. The permanent and transitory shock variances come from Sabelhaus and Song 2010. The age-education dependent mortality rates come from Brown, Liebman, and Pollet 2007.

Description	Parameter	Value
Population growth rate	N	0.0025
Technological growth rate	Γ	0.0037
Rate of high school dropouts	θ_D	0.11
Rate of high school graduates	θ_{HS}	0.55
Rate of college graduates	θ_C	0.34
Labor income tax rate	τ	0.0942

Table 3: Parameter values (annual frequency) for the life-cycle model.

The estimation yields an optimal value of $R = 1.0431$ for the R-point model in this setting. For the R-dist model in the life-cycle setting, $R = 1.002$ and $\nabla = 0.0953$. Note the improved fit to the data, as illustrated in figure 2.

Table 4 compares the relevant moments of the empirical and simulated wealth distributions again, this time for the setting with life-cycle dynamics..

Table 4: Wealth Distribution (LC)

	0–20	20–40	40–60	60–80	80–95	Top 5%
Wealth share (data)	-0.002	0.011	0.044	0.118	0.255	0.574
Wealth share (no het.)	0.018	0.056	0.110	0.207	0.321	0.287
Wealth share (het.)	0.003	0.014	0.037	0.115	0.358	0.473




Figure 2: Comparison of R-Point and R-Dist Models in the Life-Cycle Setting.

4.3 Untargeted moments

Introducing an additional source of (ex-ante) heterogeneity beyond labor income uncertainty into the representative agent framework allows the simulated distribution of wealth to better match moments of the empirical wealth distribution. But this is an expected result when allowing for an additional dimension of complexity in a model. So, although the model with heterogeneous returns does a good job of matching the given lorenz targets, we need another way to assess the model’s performance.

4.3.1 The marginal propensity to consume

The literature on heterogeneous agents consistently shows that models with a precautionary savings motive will lead to a dispersion in the marginal propensity to consume (MPC) across households. This feature is key to why the HA framework is able to produce estimates of the aggregate MPC that are close to empirical measurements, such as those measured in Jappelli and Pistaferri 2014b.⁶

In the infinite horizon setting, introducing heterogeneity changes the model’s implications for both the aggregate MPC and its dispersion across households. When all households earn the same return on their savings, the aggregate MPC is 11.3%, rising to 27.5% when heterogeneous returns are introduced. We decompose this by looking at average MPCs by wealth decile in figure 3.

As Figure 3 illustrates, MPCs are much higher when households earn different returns. When heterogeneous returns are present, the model replicates the extreme wealth concentration observed in the data. Poor households are not

⁶Jappelli and Pistaferri 2014b estimate an MPC of about 48% and document significantly larger MPCs for households with low cash-on-hand.




Figure 3: Infinite-horizon marginal propensities to consume.

just somewhat poorer—they are dramatically poorer, pushing them into the hand-to-mouth region where consumption tracks income almost one-for-one.

Because incorporating life-cycle dynamics generates a more realistic distribution of wealth (even without heterogeneous returns), MPCs are slightly higher here: an aggregate MPC of 12.6% without heterogeneity and 27.8% with heterogeneous returns (see Figure 4). However, the stark difference between the model specifications with (orange line) and without heterogeneity (blue line in Figure 4) does persist.




Figure 4: Life-cycle marginal propensities to consume.

4.3.2 Wealth shares by age cohort

Next, I include age-dependent wealth moments from the same wave of the SCF to serve as another set of untargeted moments in figure 5.

Empirical Lorenz Shares by Age (2004)

age	20th	40th	60th	80th
25-30	-0.0723	-0.0657	-0.0266	0.1099
30-40	-0.008	0.0054	0.057	0.1813
40-50	-0.0001	0.0187	0.0776	0.2178
50-60	0.0018	0.0215	0.0766	0.2126
60-70	0.0011	0.0188	0.0726	0.2081

Figure 5: Empirical Lorenz Curve Targets from the 2004 SCF.

Figures 6 and 7 present the simulated version of the untargeted moments for the model without heterogeneity 6, and with heterogeneity 7. The age-dependent Lorenz targets that arise from the model again fit the data much better when returns heterogeneity is present versus when it is absent.

Simulated Lorenz Shares by Age

age	20th	40th	60th	80th
25-30	0.0449	0.1342	0.2773	0.5253
30-40	0.0421	0.1426	0.2984	0.5325
40-50	0.0465	0.1416	0.2875	0.5116
50-60	0.0506	0.1452	0.2878	0.5069
60-70	0.0466	0.1352	0.2706	0.4837

Figure 6: Simulated Untargeted Moments without Heterogeneity (R-point).**Simulated Lorenz Shares by Age**

age	20th	40th	60th	80th
25-30	0.0229	0.0951	0.2203	0.4458
30-40	0.01	0.052	0.1423	0.3432
40-50	0.0057	0.0293	0.088	0.266
50-60	0.0044	0.0216	0.0695	0.2397
60-70	0.0039	0.0195	0.0662	0.2292

Figure 7: Simulated Untargeted Moments with Heterogeneity (R-dist).

5 Wealth versus Capital Income Tax

We've discussed the role of heterogeneous returns in generating substantial wealth inequality in the simulated model. Heterogeneity in returns also has interesting policy implications. Guvenen et al. 2023 documents that, when individuals earn different returns on their assets, a capital income tax imposes a larger burden on those more efficient with their capital. In contrast, a wealth tax imposes a larger burden on those who are unproductive with their capital.

This redistributive effect of the wealth tax leads to higher aggregate productivity, output, and ultimately larger welfare gains than the capital income tax. Following this reasoning, in this section I consider the quantitative effects of each tax scheme on the distribution of wealth when heterogeneous.

I start with an economy in which the distribution of returns matches wealth moments measured in the data. The question then becomes: how does the distribution of wealth change if a revenue equivalent tax rate is applied to every household? A capital income tax will clearly decrease wealth inequality, since only households that earn capital income will see a tax but those with zero or negative returns will see a capital tax of zero. Consequently, less wealthy households will hold a higher share of the aggregate wealth than when the capital tax is not present.

The effects of the wealth tax are less obvious, since all households hold some wealth and thus will be taxed accordingly. In this setting, the wealth tax enters the household budget constraint as follows:

$$m_{t+1} = (1 + r_t - \tau_w) k_{t+1} + \xi_{t+1}.$$

Similarly, for the capital income tax,

$$m_{t+1} = (1 + (1 - \tau_{ci}) r_t) k_{t+1} + \xi_{t+1}.$$

I compute aggregate income as GDP in this setting, and then find the wealth tax rate and the capital tax rate that would raise tax revenues equal to 1% of GDP. I follow this procedure for the estimated distribution of returns both for the infinite horizon and the life-cycle setting.

5.1 Effects on the wealth distribution

Next, I apply the tax schemes to each households and compare the resulting wealth distributions for both cases before and after each policy is implemented. Table 5 presents the results of applying each of the tax schemes in the infinite horizon 5 version of the model.

Lorenz points				
Tax scheme	20%	40%	60%	80%
None	.49%	2.4%	6.7%	16.3%
Wealth $\tau_w = .33\%$	.54%	2.8%	7.8%	18.7%
Capital income $\tau_{ci} = 4.96\%$	.59%	2.97%	7.8%	18.7%

Table 5: Tax policies in the infinite horizon setting.

Applying the tax schemes increases the share of wealth held by households at each chosen percentile, reflecting the redistributive effect of taxation on extreme wealth. Although the effect is minimal in both cases, the capital income tax more effectively reduces wealth inequality than the flat wealth tax. This result is consistent with the extant literature: when heterogeneous returns are present, a capital income tax disproportionately burdens households who use capital most productively, thereby reducing the concentration of wealth.

Next, I apply the tax schemes in the life-cycle 6 version of the model and obtain similar results: the capital income tax has a larger effect on reducing wealth inequality than the wealth tax that raises the same level of tax revenue. Table 6 presents the results, which that , the effects in the life-cycle model are less pronounced than in the infinite horizon model.

Lorenz points				
Tax scheme	20%	40%	60%	80%
None	.34%	1.8%	5.4%	16.98%
Wealth $\tau_w = .36\%$	.35%	1.8%	5.6%	17.1%
Capital income $\tau_{ci} = 7.1\%$	.37%	1.9%	5.9%	17.9%

Table 6: Tax policies in the life cycle setting.

5.2 Welfare effects of the tax policies

As an additional comparison of the two tax policies when heterogeneous returns are present, I follow Guvenen et al. 2023 and examine newborns' lifetime utility, where a newborn's return type is drawn from the estimated distribution that best fits the SCF wealth data. I summarize welfare using a consumption-equivalent measure: the percentage change in consumption under one policy that would make a newborn, before learning their return type, indifferent between the wealth tax and the capital income tax. I then decompose this comparison by estimated return type.

For this analysis, I retain my calibrated value of $\rho = 1$ such that $u(c) = \log c$. The consumption-equivalent change Δ between a wealth tax (WT) and a capital income tax (CIT) is defined by equating the lifetime utilities under each policy. Thus,

$$1 + \Delta = \exp \left(\frac{V^{CIT} - V^{WT}}{S} \right),$$

where $S = \sum_{t=0}^{\infty} (\beta \cdot \mathcal{D})^t$, and V^j is the lifetime value of a newborn under regime $j \in \{WT, CIT\}$. The parameter Δ therefore represents the *constant percentage increase in consumption each period under the wealth tax regime that would make*

a newborn as well off as under the capital income tax regime. Equivalently, if $\Delta > 0$, then the newborn prefers the capital income tax to the wealth tax.

Table 7 presents the welfare effects of the tax policies for both the infinite horizon and life-cycle settings. As a baseline for comparison, the table also includes the expected lifetime utility of the original regime with no tax policy present.

Table 7: Expected Consumption-Equivalent Welfare Changes from Tax Reform

	Infinite horizon	Life-cycle
WT vs CIT	0.20%	0.15%
WT vs No Tax	0.85%	0.35%
CIT vs No Tax	0.65%	0.20%

Notes: Entries are consumption-equivalent (CE) welfare changes, Δ , expressed as percent changes. For example, WT vs CIT is the percentage increase in consumption under the wealth tax required to make a newborn indifferent to the capital income tax. Positive values therefore indicate that the right-hand policy yields higher newborn lifetime welfare than the left-hand policy.

Results from Table 7 show that the no-tax regime is preferred to both tax policies. This is expected because the analysis is partial equilibrium and therefore abstracts from any benefit associated with government spending of the tax revenue. Relative to the no-tax benchmark, both tax instruments reduce expected lifetime utility, but the wealth tax is more costly than the capital income tax in consumption-equivalent terms. In the infinite-horizon case, for example, a newborn under the wealth tax would require roughly a 0.2% increase in consumption each period to be as well off as under the capital income tax. This distinction from Guvenen et al. 2023, where the wealth tax is preferred in welfare terms, reflects the mechanism omitted here: my framework does not allow taxes to affect aggregate productivity, output, or consumption. Accordingly, the welfare comparison in this section should be interpreted as a measure of the direct utility cost of each tax regime, not as a full general-equilibrium welfare ranking.

I next decompose these welfare effects by realized return type in the infinite-horizon setting.

The per-type results clarify why expected welfare for a newborn is higher under the capital income tax than under the wealth tax. For the lower six return types, the consumption-equivalent measure is positive, meaning that these households would require compensation under the wealth tax to be as well off as they would be under the capital income tax. The intuition is straightforward. For low-return households, and especially for those with returns below one, capital-income taxation is small or zero because there is little or no taxable capital income. A wealth tax, by contrast, lowers the gross return on assets regardless of whether the household’s return is high, low, or even below one.

Only the highest return type prefers the wealth tax in these results.

Table 8: Per-Type Consumption-Equivalent Welfare Change and Baseline Return (WT vs CIT)

	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6	Type 7
Baseline R (gross)	0.9618	0.9813	1.0009	1.0204	1.0399	1.0594	1.0790
CE Δ (WT vs CIT, %)	0.23%	0.27%	0.34%	0.32%	0.30%	0.24%	-0.31%

Notes: CE entries are consumption-equivalent welfare changes (pmv-weighted within type), expressed as percent. Positive values indicate that capital income taxation is preferred to wealth taxation for that return type. Baseline R are pre-tax gross returns by type (low $\rightarrow$ high).

The life-cycle setting produces the same qualitative pattern, shown in Table 9. Here I further decompose the effects by comparing return types within education groups.

Table 9: Per-Education Per-Type Consumption-Equivalent Welfare Change and Baseline Return (WT vs CIT)

		Type 1	Type 2	Type 3	Type 4	Type 5	Type 6	Type 7
NoHS	Baseline R (gross)	0.9200	0.9472	0.9744	1.0016	1.0289	1.0561	1.0833
	CE Δ (WT vs CIT, %)	0.118%	0.143%	0.182%	0.264%	0.340%	0.135%	-0.124%
HS	Baseline R (gross)	0.9200	0.9472	0.9744	1.0016	1.0289	1.0561	1.0833
	CE Δ (WT vs CIT, %)	0.117%	0.142%	0.184%	0.271%	0.332%	0.115%	-0.108%
College	Baseline R (gross)	0.9200	0.9472	0.9744	1.0016	1.0289	1.0561	1.0833
	CE Δ (WT vs CIT, %)	0.116%	0.142%	0.184%	0.267%	0.309%	0.098%	-0.106%

Notes: CE entries are consumption-equivalent welfare changes (pmv-weighted within type), expressed as percent. Positive values indicate that capital income taxation is preferred to wealth taxation for that return type. Baseline R are pre-tax gross returns by type (low $\rightarrow$ high) at $t = 0$.

For each education group, the consumption-equivalent measure is positive for six of the seven return types. Thus, within every education category, the lower six return types would need compensation under the wealth tax to be as well off as they would be under the capital income tax. The same economic logic applies here as in the infinite-horizon model: the capital income tax is less costly for households that are less productive with a unit of capital, whereas the wealth tax reduces returns more uniformly across the distribution.

6 Extension: Mechanism for Returns Heterogeneity

We’ve discussed the main contributions of this paper. From here on, we will discuss extensions of the model which arise naturally in the discussion of the role of heterogeneous returns in explaining wealth inequality. The first extension is regarding a potential source of returns heterogeneity in the model.

Several potential explanations have been proposed for the persistent component of returns heterogeneity. One commonly cited factor is that some indi-

viduals are business owners, and variability in their entrepreneurial talent creates greater heterogeneity in labor income than standard HA models assume. Cagetti and De Nardi 2006 and Cagetti and De Nardi 2009 are notable examples of studies that explicitly model entrepreneurial talent as a component of the consumption-saving problem for households and assess the ability of such models to match empirical wealth moments.

Although modeling differences in entrepreneurial talent as variability in labor market productivity improves the model’s ability to match wealth moments at the upper tail of the distribution, it does not adequately explain mechanism considered here. This modeling choice will result in households (firms) with high levels of wealth (capital) earning lower rates of return, and vice versa for households (firms) with low levels of wealth (capital). As shown by Fagereng et al. 2020, both average returns and the persistent, idiosyncratic components of returns increase with wealth, reflecting clear scale dependence.

Another explanation, closer to my mechanism but still distinct is financial literacy or sophistication. Lusardi and Mitchell 2014 survey models that explicitly allow households to make costly investments in financial literacy, granting them access to an investment technology offering a higher average return. Lusardi, Michaud, and Mitchell 2017 show that incorporating such endogenous financial accumulation into the standard consumption-saving framework enables models to match wealth moments particularly well, this time through the returns channel rather than the labor income channel associated with entrepreneurial talent.

As I am particularly interested in a setting with a single, safe asset available, I examine another alternative source of heterogeneity in the rate of return by drawing on literature documenting substantial variation in the banking sector, specifically the rates offered to depositors.

6.1 Deposit rates and the sensitivity of deposit levels to changes in the market interest rate

Deuffhard, Georgarakos, and Inderst 2018 among others, document substantial heterogeneity in rates of return earned by depositors. To incorporate this finding in the standard consumption-saving framework, I must identify potential sources of heterogeneity in the banking sector.

I begin by considering the balance sheet of banks and the problem they face in optimally choosing the rate to offer on deposits. In a simple setting, they accept deposits at an offered deposit rate and then hold reserves within the central banking system, earning the market interest rate. This discrepancy between the offered deposit rate and the market interest rate leaves room for banks to earn a profit.

Empirical evidence from the U.S. shows that the level of deposits at a given bank will change due to exogenous changes in the federal funds rate. In fact, Drechsler, Savov, and Schnabl 2017 propose a clear transmission channel for monetary policy: changes in the federal funds rate may lead banks to widen the interest spread they charge on deposits, which causes deposits to flee the

bank. They find a strong, negative relationship between changes in the federal funds rate and the growth rate of deposits, with a 100-basis-point increase in the federal funds rate leading to higher deposit outflows in bank branches in more concentrated markets relative to those in less concentrated markets.⁷

Thus, variation in the strength of this transmission channel across banks can be viewed as a potential source of heterogeneity in the banking sector, as certain market characteristics would make the level of deposits offered by some banks less sensitive to changes in the federal funds rate than other banks. For example, Sarkisyan and Viratyosin 2021 make a distinction between *local* and *globally integrated* banks, and show that “global banks lose much more deposits relative to local banks in response to unexpected changes in the federal funds rate”. Similarly, Adrien d’Avernas et al. 2024 document heterogeneity in deposit rates between larger and smaller banks.

Given these empirical findings, I extend the standard HA model describing household consumption-saving decisions by explicitly modeling a banking sector. In this setting, banks will each solve a similar profit-maximization problem regarding accepting deposits at an offered deposit rate and holding reserves that earn the market interest rate. The key distinction between banks in the model is how sensitive the level of deposits is to changes in the market interest rate, in line with prior empirical evidence. This analysis will help explain how heterogeneity in returns may arise for households that are otherwise the same.

6.2 Model of heterogeneous deposit rates

I consider a small, open economy in which banks and households are the optimizing agents. This analysis operates in partial equilibrium as the world interest rate is taken as given. I present a simple framework for determining banks’ optimal deposit rates. Assuming a Cobb-Douglas aggregate production function, I derive the marginal product of capital (less depreciation) consistent with the capital-to-output ratio and its empirical counterpart. This effective interest rate serves as the world or “market” interest rate and, together with the estimated distribution of returns, can be used to estimate the elasticities of foreign deposits to the deposit rates for each of the banks in the model.

6.2.1 Assumptions regarding the banking sector

In this setting, the economy features a continuum of banks, identical in all respects except for the elasticity of the level of deposits to changes in the market interest rate.

The model is static: each bank sets a deposit rate at the beginning of the horizon based solely on the market interest rate and its given elasticity, without the ability to expand its depositor base.⁸

⁷Branches in “more concentrated markets” operate in local deposit markets where a few banks hold large market shares.

⁸For example, compare a bank in a suburb area of Montana (local) versus a bank near downtown Houston, Texas (globally integrated).

Households, in turn, do not endogenously choose which banks to do business with but instead are assigned a bank at birth and remain with it until death. In this simplified setting, the banking sector merely replaces the assignment of idiosyncratic rates of returns over the time horizon in the standard model.

6.2.2 Decision problem for banks accepting deposits

The sole distinction between banks in this model is the sensitivity of their level of deposits to changes in the market interest rate, which is indexed by ε_i . This heterogeneity in deposit elasticity generates variation in returns in the model. I present the decision problem for banks using a simplified version of the framework found in Paul and Ulate 2024.

Let R^m be the market rate of return, R^d be the rate of return offered on deposits by a bank, and $S(R^d, R^m)$ be the level of deposits held at a given bank.

Banks solve

$$\max(R^m - R^d) \cdot S(R^d, R^m)$$

subject to

$$S(R^d, R^m) = A \left(\frac{R^d}{R^m} \right)^\varepsilon.$$

Importantly, the first-order condition implies that the optimal deposit rate for the i -th bank is

$$R_i^d = \frac{\varepsilon_i}{1 + \varepsilon_i} R^m.$$

This relationship plays a key role in the model: once I calibrate the model for a particular value of R^m , estimating a uniform distribution of returns using the simulated method of moments will imply a corresponding distribution of elasticities (i.e., the distribution that minimizes the distance between simulated and empirical Lorenz wealth moments). The seven discretized points therefore capture seven different deposit rates offered, resulting in varying elasticities among seven different bank types in the model. From the expression above, banks with higher values of ε_i must set R_i^d closer to R^m .

6.3 The implied distribution of bank heterogeneity

Having estimated the distribution of returns that best matches the observed wealth moments, I can use these results, along with the assumptions regarding the bank's decision problem to back out an implied distribution of ε . This parameter reflects how strongly each bank's deposit base responds to market interest rate changes. In both this simplified setting and the transmission channel empirically documented by Drechsler, Savov, and Schnabl 2017, differences in these sensitivities ultimately leads to differences in the deposit rates offered by banks.

I begin by assuming that a Cobb-Douglas aggregate production function, such that the marginal product of capital can be written as $\alpha \frac{Y}{K}$. With calibrated values of $\delta = .025$ and $\alpha = .36$, and a capital-to-output ratio 3, this setting has

an effective interest rate of $R^m = 1.095$, which can be used as the market interest rate.

Since the model with heterogeneity (i.e., the R-dist model) has seven estimated points for the uniform distribution, the implied estimated points for ε can be uniquely determined by the expression

$$\varepsilon_i = \frac{R_i^d}{R^m - R_i^d}.$$

Table 10 shows, for the infinite horizon and life-cycle settings, the estimated returns distribution that best matches 2004 SCF data on net worth and the corresponding implied elasticities.

Table 10: Estimated Returns and Implied Elasticities

Infinite Horizon		Life-cycle	
Estimated returns	Implied elasticities	Estimated returns	Implied elasticities
0.962	7.222	0.920	5.256
0.981	8.634	0.947	6.408
1.001	10.632	0.974	8.081
1.020	13.676	1.002	10.728
1.040	18.877	1.029	15.553
1.059	29.786	1.056	27.127
1.079	67.239	1.083	92.489

6.3.1 Interpreting the implied distribution of elasticities

Next, I assess the model’s performance by comparing its implications to established empirical findings on bank deposit sensitivities to changes in the federal funds rate. The implied distribution of elasticities derived from my estimation method can be directly compared to those empirical estimates to assess the validity of my model. In addition, I include wealth shares by age cohort as a set of untargeted moments to further assess the model’s fit.

By defining the deposit function as $S(\cdot) = A \left(\frac{R^d}{R^m} \right)^\varepsilon$ the parameter ε can be clearly interpreted as the elasticity of deposits to changes in the market interest rate. Formally,

$$-\varepsilon = \frac{\partial \ln S(\cdot)}{\partial \ln R^m}.$$

The elasticity parameter therefore indicates how a one-percentage-point change in the market interest rate translates into a percentage change in deposits. This formulation enables me to directly compare the implied elasticities following the estimation procedure to the empirical evidence on the transmission channel described by Drechsler, Savov, and Schnabl 2017 regarding the relationship between the federal funds rate and the level of deposits at banks. For example,

Genay and Halcomb 2004 find that a 1% change in the Fed funds rate leads to about a 3% to 4% change in the level of deposits, depending on the size of the bank.

The returns heterogeneity required to match observed wealth inequality using only safe assets (i.e., bank deposits) produces vastly overstated elasticities for the banking sector. This outcome illustrates a limitation of the model: banks must raise deposit rates to attract depositors when the federal funds rate increases, regardless of their size. Consequently, there will be less variation in the optimal deposit rates offered across the banking sector, and banks that do not offer competitive deposit rates will likely find that their depositors switch to other safe investment technologies like money market funds. Since my model does not account for the number of banks in the economy or returns heterogeneity arising from choices between safe assets, it is not too surprising that deposit rate elasticities are not well matched in this setting. That said, the ability of the model to back out a distribution of elasticities under the given assumptions is still useful.

7 Extension: Estimating a Lognormal Distribution of Returns

Using Norwegian population data, Fagereng et al. 2020 document a persistent component to individual returns, illustrated by a nonuniform distribution of estimated fixed effects in the return to net worth for individuals. Motivated by this finding, I rerun the estimation of my model under the assumption that returns are instead lognormally distributed across households. The main findings of the paper are robust to this modification.

7.0.1 Comparing the simulated wealth distributions

Since the model without heterogeneous returns is the same regardless of the distributional assumptions, in Figure 8 I show the results for the infinite horizon and life cycle models with heterogeneous returns, lognormally distributed across households.




Figure 8: Comparison of PY (left) vs LC (right) R-Dist Models.

It is clear from the figure that the results from the infinite horizon setting persist here: the life cycle setting is better at matching the observed wealth distribution than the infinite horizon setting, but they both match wealth moments especially well. Again, I include the wealth moments for each of the models and the data in Table 11.

Table 11: Wealth Distribution (Lognormal Returns)

	0–20	20–40	40–60	60–80	80–95	Top 5%
Wealth share (data)	-0.002	0.011	0.044	0.118	0.255	0.574
Wealth share (PY)	0.006	0.021	0.044	0.088	0.225	0.616
Wealth share (LC)	0.004	0.016	0.039	0.107	0.335	0.498

7.1 Untargeted moments

The aggregate MPC is 25.7% and is nearly identical in the the life-cycle setting at 26%. Figure 9 presents a breakdown of average MPCs by wealth deciles for both settings. Again, the MPC's are generally higher in the life-cycle setting due to its superior ability to match lower moments of the wealth distribution.




Figure 9: Infinite-horizon and Life-cycle marginal propensities to consume.

For the second set of untargeted moments, I include the cumulative wealth shares by age for this case of an estimated lognormal returns distribution. Again, these moments align more closely with the empirical data than the case with no heterogeneity, particularly for the wealth shares of older cohorts.

Empirical Lorenz Shares by Age (2004)

age	20th	40th	60th	80th
25-30	-0.0723	-0.0657	-0.0266	0.1099
30-40	-0.008	0.0054	0.057	0.1813
40-50	-0.0001	0.0187	0.0776	0.2178
50-60	0.0018	0.0215	0.0766	0.2126
60-70	0.0011	0.0188	0.0726	0.2081

Figure 10: Empirical Lorenz Curve Targets from the 2004 SCF.

7.2 Tax implications

I also consider the tax implications of the model in this setting. I begin by comparing the effects of each of the tax policies on wealth inequality for both the infinite horizon and the life-cycle cases. The same relationship holds here: each effect of the tax is small (since the policies are revenue-equivalent and only raise 1% of aggregate labor income), but the capital income tax is more effective at reducing wealth inequality more.

Simulated Lorenz Shares by Age

age	20th	40th	60th	80th
25-30	0.0261	0.1017	0.2285	0.4545
30-40	0.0126	0.0595	0.1531	0.3432
40-50	0.0074	0.0349	0.097	0.2547
50-60	0.0059	0.0272	0.0797	0.2296
60-70	0.0054	0.0254	0.076	0.2199

Figure 11: Simulated Untargeted Moments with Heterogeneity (R-dist).

	Lorenz points			
Tax scheme	20%	40%	60%	80%
None	0.60%	2.75%	7.27%	16.28%
Wealth $\tau_w = 0.34\%$	0.67%	3.14%	8.38%	18.86%
Capital income $\tau_{ci} = 5.08\%$	0.73%	3.37%	8.94%	20.02%

Table 12: Tax policies in the infinite horizon setting (lognormal returns across households).

	Lorenz points			
Tax scheme	20%	40%	60%	80%
None	0.42%	2.05%	5.98%	16.73%
Wealth $\tau_w = 0.33\%$	0.43%	2.12%	6.14%	16.86%
Capital income $\tau_{ci} = 4.97\%$	0.46%	2.23%	6.46%	17.72%

Table 13: Tax policies in the life-cycle setting (lognormal returns across households).

For the welfare effects in this setting, I estimate the consumption-equivalent compensation required under the wealth tax to make a newborn indifferent to the capital income tax. The conclusion is the same as in the baseline specification: before learning their type, a newborn prefers the capital income tax to the wealth tax because expected lifetime utility is higher under the capital income tax regime.

These welfare effects can again be decomposed by return type for the infinite-horizon case and by education-return type in the life-cycle scenario. The qualitative pattern is unchanged: the lower six return types prefer the capital income

Table 14: Expected Consumption-Equivalent Welfare Changes from Tax Reform (Lognormal Returns)

	Infinite horizon	Life-cycle
WT vs CIT	0.22%	0.18%
WT vs Original	0.78%	0.36%
CIT vs Original	0.56%	0.18%

Notes: Entries are consumption-equivalent (CE) welfare changes, Δ , expressed as percent changes under the assumption that individual returns follow a lognormal distribution across households. WT vs CIT is the percentage increase in consumption under the wealth tax required to make a newborn indifferent to the capital income tax. Positive values therefore indicate that the right-hand policy is preferred to the left-hand policy.

tax, while only the highest return type prefers the wealth tax. A notable difference across distributional assumptions is that the estimated lognormal distribution places fewer types below a gross return of one than the uniform distribution does. Despite this difference, the welfare ranking remains robust.

Table 15: Per-Type Consumption-Equivalent Welfare Change and Baseline Return (WT vs CIT, Lognormal Returns)

	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6	Type 7
Baseline R (gross)	0.976	1.001	1.014	1.026	1.038	1.053	1.079
CE Δ (WT vs CIT, %)	0.267%	0.342%	0.329%	0.317%	0.308%	0.304%	-0.326%

Notes: CE entries are per-type consumption-equivalent welfare changes (pmv-weighted within type), expressed as percent. Positive values indicate that the capital income tax is preferred to the wealth tax for that return type. Baseline R values are pre-tax gross returns (low $\rightarrow$ high) under the lognormal returns specification.

Table 16: Per-Education Per-Type Consumption-Equivalent Welfare Change and Baseline Return (WT vs CIT, Lognormal Returns)

		Type 1	Type 2	Type 3	Type 4	Type 5	Type 6	Type 7
NoHS	Baseline R (gross)	0.9363	0.9711	0.9910	1.0084	1.0260	1.0471	1.0865
	CE Δ (WT vs CIT, %)	0.132%	0.176%	0.225%	0.275%	0.335%	0.292%	-0.114%
HS	Baseline R (gross)	0.9363	0.9711	0.9910	1.0084	1.0260	1.0471	1.0865
	CE Δ (WT vs CIT, %)	0.131%	0.177%	0.229%	0.283%	0.332%	0.250%	-0.098%
College	Baseline R (gross)	0.9363	0.9711	0.9910	1.0084	1.0260	1.0471	1.0865
	CE Δ (WT vs CIT, %)	0.131%	0.177%	0.228%	0.276%	0.311%	0.218%	-0.096%

Notes: CE entries are consumption-equivalent welfare changes (pmv-weighted within type), expressed as percent. Positive values indicate that capital income taxation is preferred to wealth taxation for that return type. Baseline R are pre-tax gross returns by type (low $\rightarrow$ high) at $t = 0$ under the lognormal returns specification.

7.3 Mechanism for bank heterogeneity

Table 17 captures the estimated returns distribution that best matches the 2004 SCF data on net worth and the corresponding implied elasticities for both the infinite horizon and life-cycle settings.

Table 17: Estimated Returns and Implied Elasticities (Lognormal)

Infinite Horizon		Life-cycle	
Estimated returns	Implied elasticities	Estimated returns	Implied elasticities
0.976	8.231	0.936	5.898
1.001	10.615	0.971	7.838
1.014	12.602	0.991	9.532
1.026	14.963	1.008	11.640
1.038	18.366	1.026	14.874
1.053	24.922	1.047	21.857
1.079	68.242	1.086	127.378

8 Conclusion

This paper explores the ability of macroeconomic models to generate a distribution of wealth with substantial inequality by the estimation of a calibrated consumption-saving model that allows for heterogeneous returns. Consistent with other deviations from the representative agent framework, I find that the ex-ante heterogeneity in rates of return needed to match wealth moments compares well with empirical estimates of the returns to net worth found by recent work.

Unlike much of the existing literature linking persistent return heterogeneity to factors such as entrepreneurial ability or financial sophistication, I focus more on heterogeneity in deposit rates across banks. I incorporate related literature in the standard HA framework under a simple but realistic assumption that many households remain “stuck” with the bank in or around their neighborhood. That bank makes complex financial decisions that ultimately affect households through the channel of varying deposit rates offered.

Although I do not allow households to switch banks, this mechanism is similar to financial literacy explanations for returns heterogeneity, without the need to model portfolio risk. This exclusion is useful because (i) untangling how much of the persistent component of returns comes from risk preferences and from financial sophistication is not straightforward and (ii) there is significant heterogeneity in returns even when individuals hold no risky assets.

This model is a partial equilibrium analysis. The market interest rate is being taken as given. It is not determined by some market clearing condition. I view my model as the simplest implementation of a potential source of heterogeneity. In the simulation of the model and the resulting SMM estimation, I do not add banks as an agent type. Thus, the bank is not responding in every period to the

level of deposits they receive after they set the optimal deposit rate based on the demand for deposits. Consequently, I avoid having to choose a particular scheme of allocating agents in the model to a particular bank. In this way, there are seven types of banks just as there are seven types of returns that an agent may receive.

The culmination of these simplifying assumptions leaves us with a setting where we can consider the most stark role for returns in explaining wealth inequality. With more features in the model, like general equilibrium considerations, endogenous returns heterogeneity, portfolio choice, overlapping generations and bequests, etc., we leave more room for these features to explain the generated skewness in the model’s distribution of wealth.

References

- Adrien d’Avernas et al. (Aug. 2024). *The Deposit Business at Large vs. Small Banks*. URL: <https://www.fdic.gov/system/files/2024-09/wallace-paper-091224.pdf>.
- Aiyagari, S Rao (1994). “Uninsured idiosyncratic risk and aggregate saving”. In: *The Quarterly Journal of Economics* 109.3, pp. 659–684.
- Altmejd, Adam, Thomas Jansson, and Yigitcan Karabulut (2024). *Business education and portfolio returns*. IZA-Institute of Labor Economics.
- Bach, Laurent, Laurent E. Calvet, and Paolo Sodini (2018). “Rich Pickings? Risk, Return, and Skill in Household Wealth”. In: *American Economic Review* 110.9, pp. 2703–47. DOI: 10.1257/aer.20170666. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20170666>.
- Benhabib, Jess and Alberto Bisin (2018). “Skewed Wealth Distributions: Theory and Empirics”. In: *Journal of Economic Literature* 56.4, pp. 1261–91. DOI: 10.1257/jel.20161390. URL: <https://www.aeaweb.org/articles?id=10.1257/jel.20161390>.
- Benhabib, Jess, Alberto Bisin, and Mi Luo (2017). “Earnings Inequality and Other Determinants of Wealth Inequality”. In: *American Economic Review* 107.5, pp. 593–97. DOI: 10.1257/aer.p20171005. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.p20171005>.
- (May 2019). “Wealth Distribution and Social Mobility in the US: A Quantitative Approach”. In: *Am. Econ. Rev.* 109.5, pp. 1623–1647. ISSN: 0002-8282. DOI: 10.1257/aer.20151684. URL: <https://www.aeaweb.org/doi/10.1257/aer.20151684>.
- Bewley, Truman (1983). “A difficulty with the optimum quantity of money”. In: *Econometrica: Journal of the Econometric Society*, pp. 1485–1504.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008). “Consumption Inequality and Partial Insurance”. In: *American Economic Review* 98.5, pp. 1887–1921. DOI: 10.1257/aer.98.5.1887.
- Brown, Jeffrey R, Jeffrey B Liebman, and Joshua Pollet (Nov. 2007). “Appendix: Estimating Life Tables That Reflect Socioeconomic Differences in Mortality”. en. In: *The Distributional Aspects of Social Security and Social Security Re-*

- form*. University of Chicago Press, pp. 447–458. ISBN: 9780226241890. URL: <https://www.degruyter.com/document/doi/10.7208/9780226241890-013/html?lang=en>.
- Cagetti, Marco (2003). “Wealth Accumulation over the Life Cycle and Precautionary Savings”. In: *J. Bus. Econ. Stat.* 21.3, pp. 339–353. ISSN: 0735-0015. URL: <http://www.jstor.org/stable/1392584>.
- Cagetti, Marco and Mariacristina De Nardi (Oct. 2006). “Entrepreneurship, Frictions, and Wealth”. In: *J. Polit. Econ.* 114.5, pp. 835–870. ISSN: 0022-3808,1537-534X. DOI: 10.1086/508032. URL: <https://doi.org/10.1086/508032>.
- (Mar. 2009). “Estate Taxation, Entrepreneurship, and Wealth”. In: *Am. Econ. Rev.* 99.1, pp. 85–111. ISSN: 0002-8282. DOI: 10.1257/aer.99.1.85. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.99.1.85>.
- Campbell, John Y, Tarun Ramadorai, and Benjamin Ranish (Sept. 2019). “Do the Rich Get Richer in the Stock Market? Evidence from India”. In: *American Economic Review: Insights* 1.2, pp. 225–240. DOI: 10.1257/aeri.20180158. URL: <https://www.aeaweb.org/articles?id=10.1257/aeri.20180158>.
- Carroll, Christopher et al. (2017). “The distribution of wealth and the marginal propensity to consume”. In: *Quantitative Economics* 8.3, pp. 977–1020. DOI: <https://doi.org/10.3982/QE694>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/QE694>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.3982/QE694>.
- Carroll, Christopher D (1992). “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence”. In: *Brookings Pap. Econ. Act.* 1992.2, pp. 61–156. ISSN: 0007-2303.
- Carroll, Christopher D, Jiri Slacalek, and Kiichi Tokuoka (July 2015). “Buffer-stock saving in a Krusell-Smith world”. In: *Econ. Lett.* 132, pp. 97–100. ISSN: 0165-1765. DOI: 10.1016/j.econlet.2015.04.021. URL: <https://www.sciencedirect.com/science/article/pii/S016517651500172X>.
- Daminato, Claudio and Luigi Pistaferri (May 2024). “Returns Heterogeneity and Consumption Inequality Over the Life Cycle”. DOI: 10.3386/w32490. URL: <http://www.nber.org/papers/w32490>.
- De Nardi, Mariacristina and Giulio Fella (Oct. 2017). “Saving and wealth inequality”. In: *Rev. Econ. Dyn.* 26, pp. 280–300. ISSN: 1094-2025. DOI: 10.1016/j.red.2017.06.002. URL: <https://www.sciencedirect.com/science/article/pii/S1094202517300546>.
- Debacker, Jason et al. (2013). “Rising Inequality: Transitory or Persistent? New Evidence from a Panel of U.S. Tax Returns”. In: *Brookings Pap. Econ. Act.*, pp. 67–122. ISSN: 0007-2303, 1533-4465. URL: <http://www.jstor.org/stable/23594863>.
- Den Haan, Wouter J, Kenneth L Judd, and Michel Juillard (Jan. 2010). “Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty”. In: *J. Econ. Dyn. Control* 34.1, pp. 1–3. ISSN: 0165-1889. DOI: 10.1016/j.jedc.2009.07.001. URL: <https://www.sciencedirect.com/science/article/pii/S0165188909001286>.

- Deuflhard, Florian, Dimitris Georgarakos, and Roman Inderst (Apr. 2018). “Financial Literacy and Savings Account Returns”. en. In: *J. Eur. Econ. Assoc.* 17.1, pp. 131–164. ISSN: 1542-4766. DOI: 10.1093/jeea/jvy003. URL: <https://academic.oup.com/jeea/article-abstract/17/1/131/4981453>.
- Drechsler, Itamar, Alexi Savov, and Philipp Schnabl (Nov. 2017). “The deposits channel of monetary policy”. en. In: *Q. J. Econ.* 132.4, pp. 1819–1876. ISSN: 0033-5533, 1531-4650. DOI: 10.1093/qje/qjx019. URL: <https://dx.doi.org/10.1093/qje/qjx019>.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik (2021). “MPC Heterogeneity and Household Balance Sheets”. In: *American Economic Journal: Macroeconomics* 13.4, pp. 1–54. DOI: 10.1257/mac.20190211.
- Fagereng, Andreas et al. (2020). “Heterogeneity and Persistence in Returns to Wealth”. In: *Econometrica* 88.1, pp. 115–170. DOI: <https://doi.org/10.3982/ECTA14835>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA14835>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA14835>.
- Friedman, Milton (1957). *Theory of the Consumption Function*. Princeton University Press. URL: <http://www.jstor.org/stable/j.ctv39x7zh> (visited on 03/21/2023).
- Gabaix, X et al. (2016). “The dynamics of inequality”. In: *Econometrica*. ISSN: 0012-9682. DOI: 10.3982/ECTA13569. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13569?casa_token=9ibahWJ9lG4AAAAA:OdYLGnKVIEPFH4Qjz0khrmrwUun5w7L47DXcniNeY0AnQM5FIFHG5164k1_OLkoZNmHKtLN4v9HI5TA.
- Genay, Hesna and Darrin R Halcomb (Nov. 2004). *Rising Interest Rates, Bank Loans and Deposits - Federal Reserve Bank of Chicago*. en. <https://www.chicagofed.org/publications/chicago-fed-letter/2004/november-208>. Accessed: 2025-8-4.
- Guler, Bulent, Burhan Kuruscu, and Baxter Robinson (2022). *The composition and distribution of wealth and aggregate consumption dynamics*. <https://events.bse.eu/live/files/4096-gkrdraftv31submitpdf>. Accessed: 2023-9-6. URL: <https://events.bse.eu/live/files/4096-gkrdraftv31submitpdf>.
- Guvenen, Fatih et al. (Apr. 2023). “Use It or Lose It: Efficiency and Redistributive Effects of Wealth Taxation”. In: *Q. J. Econ.* 138.2, pp. 835–894. ISSN: 0033-5533. DOI: 10.1093/qje/qjac047. URL: <https://academic.oup.com/qje/article-pdf/138/2/835/49730065/qjac047.pdf>.
- Havránek, Tomáš and Anna Sokolova (2020). “Do Consumers Really Follow a Rule of Thumb? Three Thousand Estimates from 144 Studies Say “Probably Not””. In: *Review of Economic Dynamics* 35, pp. 97–122. DOI: 10.1016/j.red.2019.05.004.
- Huggett, Mark (1993). “The risk-free rate in heterogeneous-agent incomplete-insurance economies”. In: *Journal of economic Dynamics and Control* 17.5-6, pp. 953–969.
- Jappelli, Tullio and Luigi Pistaferri (2014a). “Fiscal Policy and MPC Heterogeneity”. In: *American Economic Journal: Macroeconomics* 6.4, pp. 107–136. DOI: 10.1257/mac.6.4.107.

- Jappelli, Tullio and Luigi Pistaferri (2014b). “Fiscal Policy and MPC Heterogeneity”. In: *American Economic Journal: Macroeconomics* 6.4, pp. 107–36. DOI: 10.1257/mac.6.4.107. URL: <https://www.aeaweb.org/articles?id=10.1257/mac.6.4.107>.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles (2006). “Household Expenditure and the Income Tax Rebates of 2001”. In: *American Economic Review* 96.5, pp. 1589–1610. DOI: 10.1257/aer.96.5.1589.
- Kaplan, Greg and Giovanni L Violante (2022). *The Marginal Propensity to Consume in Heterogeneous Agent Models*. Working Paper 30013. National Bureau of Economic Research. DOI: 10.3386/w30013. URL: <http://www.nber.org/papers/w30013>.
- Krusell, Per and Anthony Smith (1998). “Income and Wealth Heterogeneity in the Macroeconomy”. In: *Journal of Political Economy* 106.5, pp. 867–896. DOI: 10.1086/250034. URL: <https://ideas.repec.org/a/ucp/jpolec/v106y1998i5p867-896.html>.
- Lee, Seungcheol, Ralph Luetticke, and Morten O. Ravn (Dec. 2021). *Financial frictions: micro vs macro volatility*. ECB Working Paper Series 2622. European Central Bank. DOI: 10.2866/13857. URL: <https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2622~cf6025d119.en.pdf>.
- Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S. Mitchell (2017). “Optimal Financial Knowledge and Wealth Inequality”. In: *Journal of Political Economy* 125.2, pp. 431–477. DOI: 10.1086/690950. eprint: <https://doi.org/10.1086/690950>. URL: <https://doi.org/10.1086/690950>.
- Lusardi, Annamaria and Olivia S Mitchell (Mar. 2014). “The Economic Importance of Financial Literacy: Theory and Evidence”. en. In: *J. Econ. Lit.* 52.1, pp. 5–44. ISSN: 0022-0515. DOI: 10.1257/jel.52.1.5. URL: <http://dx.doi.org/10.1257/jel.52.1.5>.
- Menzio, Guido and Saverio Spinella (May 2025). *A Quantitative Theory of Heterogeneous Returns to Wealth*. URL: <https://bpb-us-e1.wpmucdn.com/wp.nyu.edu/dist/e/11962/files/2025/05/KBJ.pdf>.
- Parker, Jonathan A. et al. (2013). “Consumer Spending and the Economic Stimulus Payments of 2008”. In: *American Economic Review* 103.6, pp. 2530–2553. DOI: 10.1257/aer.103.6.2530.
- Paul, Pascal and Mauricio Ulate (Apr. 2024). “A macroeconomic model of central bank digital currency”. In: *Federal Reserve Bank of San Francisco, Working Paper Series* 2024.11, pp. 01–83. DOI: 10.24148/wp2024-11. URL: <https://www.frbsf.org/research-and-insights/publications/working-papers/2024/04/macroeconomic-model-of-central-bank-digital-currency/>.
- Sabelhaus, John and Jae Song (May 2010). “The great moderation in micro labor earnings”. In: *J. Monet. Econ.* 57.4, pp. 391–403. ISSN: 0304-3932. DOI: 10.1016/j.jmoneco.2010.04.003. URL: <https://www.sciencedirect.com/science/article/pii/S0304393210000358>.
- Sarkisyan, Sergey and Tasaneeya Viratyosin (2021). “The impact of the deposit channel on the international transmission of monetary shocks”. en. In: *SSRN*

Electron. J. ISSN: 1556-5068. DOI: 10.2139/ssrn.3938284. URL: <http://dx.doi.org/10.2139/ssrn.3938284>.

Suen, Richard M. H. (Jan. 2012). *Time Preference and the Distributions of Wealth and Income*. Working Paper 2012-01. University of Connecticut. URL: <https://media.economics.uconn.edu/working/2012-01.pdf>.