Monetary Policy, Heterogeneous Returns, and Top Inequality*

Jacob Sundram†

March 11, 2026

Abstract

Who are the winners and losers from monetary policy? This is a key question in macroeconomics. Yet, state-of-the-art models miss two crucial aspects of the data that shape the answer: Large wealth inequality and strong pass-through of aggregate returns to the rich. I add heterogeneous returns to a Heterogeneous Agent New Keynesian (HANK) model to match both. The rich gain disproportionately from easy monetary policy in this model, with the top 0.1% gaining 11% of the income from monetary policy—an order of magnitude more than in standard models. The difference arises because the rich are wealthier due to earning higher returns, and because capital income created by lower interest rates benefits them more due to the heterogeneous pass-through of returns. Instead, fiscal policy is more equal. Thus, poor and rich households disagree about policy: Poor households prefer fiscal expansion to stabilize the economy in recessions and monetary tightening in booms, while the rich prefer the opposite. Policymakers concerned about inequality should consider this.

Keywords: Monetary policy, business cycles, heterogeneous households
JEL Codes: D31, E21, E32, E52

*I am thankful for useful comments from Adrien Auclert, Jeppe Druedahl, Brigitte Hochmuth, Raphaël Huleux, John Kramer, Moritz Lenel, Manuel Menkhoff, Morten Ravn, Søren Hove Ravn, Karthik Sastry, Gianluca Violante, and seminar participants at the 2025 Annual Meeting Society for Economic Dynamics, the 17th Nordic Macro Summer Symposium, the 2025 DAEiNA Meeting, the University of Copenhagen, and Princeton University. I am grateful for financial support from the Carlsberg Foundation (Grant CF20-0546). First version: April 2025.

†Department of Economics, University of Copenhagen, jacobmsundram@gmail.com

1 Introduction

Who are the winners and losers from monetary policy? This question has received increasing interest in the last decade, both because it can shape the aggregate effects of monetary policy and also because it is important if policymakers care about inequality. This highlights the need for models that take the interaction of inequality and monetary policy seriously. Heterogeneous Agent New Keynesian (HANK) models do this by adding heterogeneity to the New Keynesian model with nominal rigidities.[1]

However, these models miss two aspects of heterogeneity that shape monetary policy's distributional effects. First, they severely understate top wealth inequality. In benchmark HANK models, no one has more than $10 million in wealth despite these households owning almost 1/3 of US wealth. Second, they assume that everyone earns the same return on their wealth. I show that this is at odds with the data by building a panel dataset of returns for US households in the Panel Study of Income Dynamics. Here, (i) returns are dispersed, (ii) the rich earn higher returns, and (iii) the returns of the rich co-move more with the aggregate. I show that both shortcomings matter because they shape sufficient statistics for monetary policy's effects on income.

For this reason, I build a HANK model with heterogeneous returns and calibrate it to the data on households' returns. I have three main takeaways. First, heterogeneous returns significantly improve the micro fit of the model. In particular, the model matches top wealth inequality, as some households get persistently high returns on their wealth. Additionally, the returns of the rich co-move more with the aggregate.

This shapes the second result: Monetary policy mostly benefits the rich. To see why, consider a central bank that cuts the interest rate. This increases capital income and thus returns as the interest rate is used to discount profits. With heterogeneous returns, this mainly benefits the rich as (i) they hold more assets and (ii) the pass-through of returns to them is stronger. In particular, the top 0.1% gain 11% of income due to monetary policy—an order of magnitude more than in standard models.

Third, households disagree about policies. In particular, consider a policymaker who wants to stimulate the economy in a recession. The rich like monetary policy, while the poor favor fiscal expansion as it mostly boosts labor income, which is more equal. However, this reverses in booms. Thus, both favor an asymmetric approach over the business cycle, using one policy in recessions and the other in booms.

1. Kaplan et al. (2018), Auclert et al. (2024b), Bayer et al. (2024), and many more.

The key to the results is heterogeneous returns. Heterogeneous returns have received significant interest in the household finance literature, but this paper is the first to add them to a HANK model.[2] I add heterogeneous returns to the model by letting households' returns follow a Markov chain that averages to the aggregate return in the economy. I then let the pass-through of aggregate to households' returns be heterogeneous and depend on wealth. This is often done for labor earnings instead of returns on wealth in the HANK literature, so my approach is a generalization.

I calibrate the returns to a panel dataset of returns across US households from 2000 to 2018 in the PSID. I use the data to establish a new empirical fact: The co-movement over time between the average return and households' return is increasing in the households' wealth. This means that in periods when returns are high, returns for rich households are even higher, while returns for poor households tend to stay the same. Numerically, the returns of the rich rise almost twice as much as average returns, while the returns of the poorest do not rise at all.

In addition to this new fact, the data also reveals strong heterogeneity in returns across the wealth distribution. In particular, rich households earn higher returns than poor households: The bottom 20% by wealth earn returns of -1% per year on average, while the top 20% earn around 2.5%. In addition to the PSID, I corroborate my results using two alternative datasets: The Survey of Consumer Finances and the Distributional National Accounts.

Using my HANK model calibrated to match the heterogeneous returns in the data, I show three main results. The first result is that adding heterogeneous returns to a standard HANK model improves the fit of the model to the microeconomic data. This is important not just for empirical realism, but because it matters for the winners and losers of monetary policy. In particular, I show that four variables shape the distribution of income in response to monetary policy. These are sufficient statistics in the sense that any model that matches these will yield the same distribution of income in response to monetary policy. The components of the sufficient statistics are: (i) The wealth distribution, (ii) the income distribution, (iii) the pass-through of aggregate returns to households' returns, and (iv) the pass-through of aggregate earnings to households' earnings. This motivates that the model should not just match the cross-sectional dispersion in returns, but also the heterogeneous pass-through of returns over the business cycle. The model with heterogeneous returns fits

2. See for instance Fagereng et al. (2020) and Bach et al. (2020).

all four components. Crucially, standard HANK models struggle with fitting these components. In particular, they feature much too little wealth inequality at the top and a common pass-through of aggregate returns to households' returns.

The key mechanism that makes the model with heterogeneous returns fit the microeconomic data is that returns are increasing in wealth. This occurs endogenously in the model with persistent heterogeneous returns: Households who earn high returns are wealthier both directly because they get more capital income, but also because they *choose* to save more as they know the return will tend to be high next period. This is a multiplying effect that makes some households very rich, so the model replicates the concentration of wealth at the top. This is a well-known fact that remains elusive to standard HANK models, which understate the wealth of the richest households by several orders of magnitude. The model with heterogeneous returns replicates the wealth concentration despite featuring no permanent heterogeneity, no preference heterogeneity, and only a single asset.

The model with heterogeneous returns also matches a key moment in business cycle analysis: A high marginal propensity to consume (MPC). The MPC has been emphasized in the literature as a key moment when studying business cycles due to its role in shaping the Keynesian multiplier. Despite this, standard HANK models with common returns struggle with matching a realistically high MPC without other additions to the model. In particular, these models face an MPC-wealth trade-off (Kaplan and Violante 2022): Replicate a high MPC but have a wealth level an order of magnitude lower than in the data, or replicate the wealth in the data but have an MPC well outside the range of empirical estimates. A model with heterogeneous returns simultaneously matches both: Some households earn high returns and are very rich, increasing the average wealth. Other households earn low returns and are close to the borrowing constraint, increasing the average MPC.

Having shown the micro fit of the model, I turn to the second result: The benefits of monetary policy mostly accrue to the rich. To show this, I ask: *For each $100 generated by monetary policy, how much goes to the top x%?* I find that the richest 0.1% gain 11% of the total increase in income on impact when monetary policy is eased. In the standard HANK model, the top 0.1% gain less than 2% from monetary policy. On the other hand, fiscal policy is similar to standard models and much more equal.

The rich gain so much when returns are heterogeneous because of the interaction between the concentration of wealth and capital income among the rich and the increase in capital income induced by expansionary monetary policy. Expansionary monetary policy increases capital income for two reasons: Directly through lower

discount rates on firm profits and indirectly through higher profits, which are procyclical. In other words, a lower *interest rate* implies higher *returns* on wealth on impact. The increase in capital income almost entirely goes to the top of the wealth distribution. In the standard HANK model, both top wealth and capital income at the top are significantly understated. Thus, the capital income gains at the top are missing or severely understated.

While the effects on income are intuitive, what households ultimately care about is welfare. For this reason, I compute the equivalent variations of monetary policy for different households. This answers the question: *How many dollars should a household be given to be indifferent between facing the monetary policy shock and not facing it?* This welfare measure takes into account the full dynamic path of the shock. The results from this exercise are very similar compared to when looking at income. In particular, rich households prefer monetary policy while poor households prefer fiscal policy.

Having studied the unequal effects of policies, I turn to the third result: Poor households prefer stabilizing recessions with fiscal policy and booms with monetary policy, and vice versa for the rich. To elaborate on this, consider a policymaker who wants to stabilize business cycles. If the gains of monetary policy are very uneven, while fiscal policy is more evenly distributed, one might be tempted to conclude that poor households prefer using fiscal policy to stabilize business cycles. However, the argument is symmetric over the business cycle: Poor households dislike tightening fiscal policy in booms just as much as they like easing fiscal policy in recessions. Likewise, for rich households with monetary policy.

For this reason, I study asymmetric policies over the business cycle. In particular, I estimate three shock series such that the model generates business cycles like those in the US. I then simulate the economy facing these shocks under two policy regimes: One using fiscal policy to stabilize the economy in recessions and monetary policy in booms, and one doing the opposite. I find the poor prefer the first policy, while the rich favor the second. This shows that both groups prefer asymmetric policy over the business cycle, but they prefer different ones. This makes clear the advantages of asymmetric policy, but leaves open the optimal policy as a normative question weighing the preferences of different groups of households.

Ultimately, this highlights the importance of taking seriously the fit to microeconomic data in macroeconomics. A policymaker who cares not only about aggregate stabilization but also who gains from different policies should consider this.

1.1 Related Literature

My paper contributes to three strands of the literature. First, my paper contributes to the literature on HANK models. A seminal paper in this literature is Kaplan et al. (2018), who introduce the concept of a HANK model and study the transmission of monetary policy. This is expanded further upon in Kaplan and Violante (2018) and Alves et al. (2020). Several other papers study the transmission of monetary policy in other HANK models, including Broer et al. (2020) and Auclert et al. (2020). Fiscal policy has also received a lot of interest in the HANK literature, with contributions counting Hagedorn et al. (2019), Auclert et al. (2023), and Auclert et al. (2024a).

Other papers study aspects such as the transmission of foreign shocks (de Ferra et al. 2020), emphasize the importance of labor markets (Ravn and Sterk 2021), or estimate a HANK model (Bayer et al. 2024). Common to all these papers is that they assume common returns.[3] On the other hand, my key contribution here is to add heterogeneous returns to the model, which allows me to take the distributional effects of shocks more seriously due to replicating the microeconomic data. Lastly, an early key paper in the HANK literature is Gornemann et al. (2016). This paper is close to mine in studying the distributional effects of monetary policy. The key differences of my paper compared to theirs are adding heterogeneous returns and comparing monetary to fiscal policy instead of comparing monetary policy hawkishness.

Second, my paper contributes to the literature on the distributional effects of shocks and, in particular, monetary policy. Key contributions are Coibion et al. (2017), Holm et al. (2021), Andersen et al. (2023), and McKay and Wolf (2023a). My contribution to the theoretical literature is to replicate the empirical distributions, which is key to understanding the *distributional* effects of shocks. My contribution compared to the empirical literature is to study welfare and policy. Additionally, my contribution is to emphasize the importance of the very top of the wealth distribution, say the top 0.1%.

Third, my paper contributes to the literature on heterogeneous returns. Heterogeneous returns are often used in household finance. The literature on heterogeneous returns can be split into two: Theory and evidence. Key theoretical contributions include Benhabib et al. (2011), Benhabib et al. (2015), Gabaix et al. (2016), Benhabib et al. (2017), Jones and Kim (2018), and Guvenen et al. (2023). These papers show how heterogeneous returns help explain wealth concentration and dynamics. There have also been a number of empirical contributions building datasets of heteroge-

3. At least within an asset class.

neous returns. This literature includes papers based on register data such as Bach et al. (2020) for Sweden, Fagereng et al. (2020) for Norway, and Smith et al. (2022) for the US. Other papers use survey data, such as Xavier (2021), who uses the Survey of Consumer Finances, and Snudden (2021) and Daminato and Pistaferri (2024), who use the PSID. My contribution is to embed heterogeneous returns in a HANK model and to study the implications for business cycle dynamics and welfare.[4]

2 Empirics

In this section, I present two key aspects of the microeconomic data that standard Heterogeneous Agent New Keynesian (HANK) models miss: The concentration of wealth at the top and the heterogeneity in returns. The goal of doing so is to build a HANK model that replicates these aspects and then study the distributional aspects of monetary policy in this model.

2.1 Concentration of Wealth

I start by discussing the concentration of wealth at the top of the wealth distribution. It is well known that wealth is highly concentrated among the wealthiest households. For instance, the top 0.1% of households hold 14% of all wealth in the US. HANK models in the literature usually do not match this, as I show in Appendix A.1.

Furthermore, data shows that the wealth distribution is "fat-tailed". In particular, it has been shown that the right tail of the wealth distribution is well approximated by a Pareto distribution, which has a fat right tail. Mathematically, I note that if a variable X is Pareto distributed, it holds that

$$\log P(X > x) \sim -\alpha \log x \quad \text{as} \quad x \to \infty, \tag{1}$$

where α is the Pareto tail index and $\log P(X > x)$ is the log counter-CDF (CCDF). The tail index then controls how "fat" the tail is, i.e., how concentrated wealth is at the top. In particular, a lower tail index, α, corresponds to *more* concentrated

4. Another paper that studies the interaction of monetary policy and heterogeneous returns is Menzio and Spinella (2025), which is contemporaneous with my paper. Compared to the literature and my paper, their paper is about *what* microeconomic foundations generate the dispersion in returns. In contrast with my paper, they do not fit the top of the wealth distribution and do not emphasize MPCs. Additionally, their model does not feature nominal rigidities and is therefore not a HANK model.

wealth. An estimate of the Pareto tail index in the US is 1.52. This implies a significant concentration of wealth at the top. For instance, this means that the variance of wealth is undefined. HANK models in the literature usually are not fat-tailed and therefore do not have a tail index, cf. Appendix A.1.

To get at this visually, Figure 1 shows the Pareto tail in the US data. The figure also shows the right tail in two HANK models from influential papers in the literature (Kaplan et al. 2018 and Auclert et al. 2024b).[5] I first note that the HANK model of Auclert et al. (2024b) provides a poor fit to the right tail of the wealth distribution. Consider next the two-asset model of Kaplan et al. (2018), which provides the best fit to top wealth shares of the models considered in Appendix A.1. Despite providing a better fit to the top than the other models in the literature, this model also does not match the right tail. In particular, the richest households in these models have little more than 10 million USD in wealth, while the richest households in the data have several orders of magnitude more.

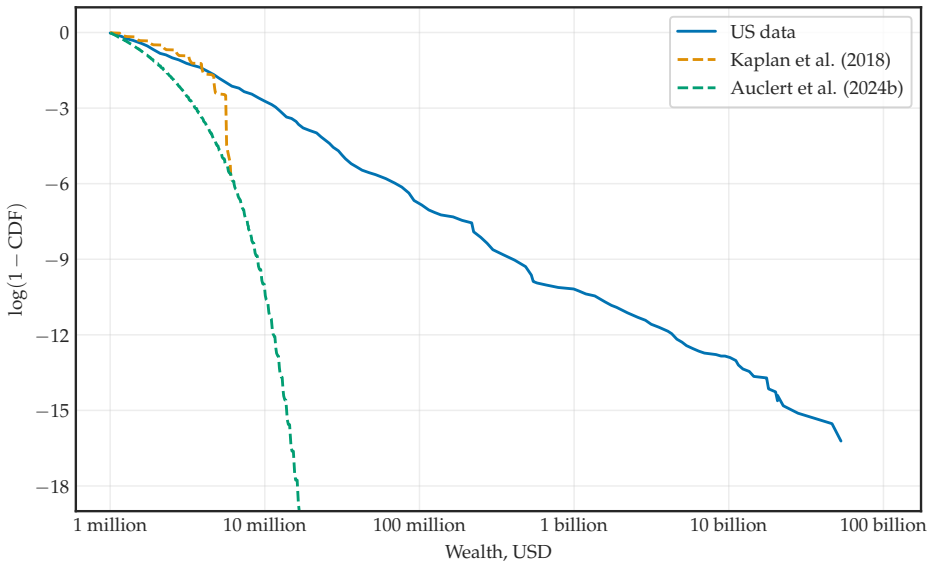


Figure 1: Wealth Concentration at the Top

Note: The figure shows the distribution of wealth in the model from Kaplan et al. (2018) and the "HA-hi-liq" model from Auclert et al. (2024b). It does so by plotting the counter-CDF against the level of wealth in USD. Both axes are log-scale. The figure also shows the counter-CDF for the US based on data from Guvenen et al. (2023) and Vermeulen (2018).

5. These papers were chosen based on being highly influential in the HANK literature, being calibrated to a high level of wealth, and having replication packages publicly available.

2.2 Return Heterogeneity

Having studied the concentration of wealth at the top, I now turn to the degree of return heterogeneity. To study heterogeneous returns, I take two different approaches using different datasets. First, I study return heterogeneity in two datasets: The Survey of Consumer Finances (SCF) and the Distributional National Accounts due to Piketty et al. (2018) (PSZ). In both datasets, I compare the joint distributions of capital income and wealth. Second, I directly construct a dataset of heterogeneous returns across US households using the Panel Study of Income Dynamics (PSID). The two approaches complement each other. The advantage of the first approach is that it is straightforward, using readily available data, and is less prone to measurement error. The advantage of the second approach is that it gets directly at the heterogeneous returns, which allows me to study additional aspects.

2.2.1 Capital Income in the SCF and PSZ

I start by studying heterogeneous returns in the cross-section. I do this using two datasets: The SCF and the Distributional National Accounts due to PSZ. The details of the data are provided in Appendices A.2 and A.3. This approach requires neither measuring returns directly nor a panel. The approach starts by computing shares of wealth and capital income. To do so, note that capital income is defined by

$$x_{it} = r^a_{it} a_{it-1}.$$

If returns are common across households, $r^a_{it} = r^a_t$, capital income can be written as $x_{it} = r^a_t a_{it-1}$. Consider now looking at the bottom p% of households in the wealth distribution. Denote these households by $i \in \mathcal{P}$. How large a share of total wealth and capital income is held by these households? In the case of common returns, the bottom shares of wealth and capital income are given by (neglecting time subscripts):

$$S(a) \equiv \frac{\sum_{i\in\mathcal{P}} a_i}{\sum a_i} = \frac{r^a \sum_{i\in\mathcal{P}} a_i}{r^a \sum a_i} = \frac{\sum_{i\in\mathcal{P}} x_i}{\sum x_i} \equiv S(x).$$

Intuitively, if returns are common, the share of wealth and capital income held by the bottom p% is the same. If returns are heterogeneous, they can be different. This allows me to test if returns are heterogeneous. Figure 2 plots this. In particular, the figure shows the share of wealth held by the bottom p% in the wealth distribution and their share of capital income. The figure clearly shows the shares of capital income

as a function of the shares of wealth lying below the 45-degree line, i.e., $S(x) < S(a)$. This implies that returns are heterogeneous. Not only this, it implies that wealthier households earn higher *rates* of return on average.

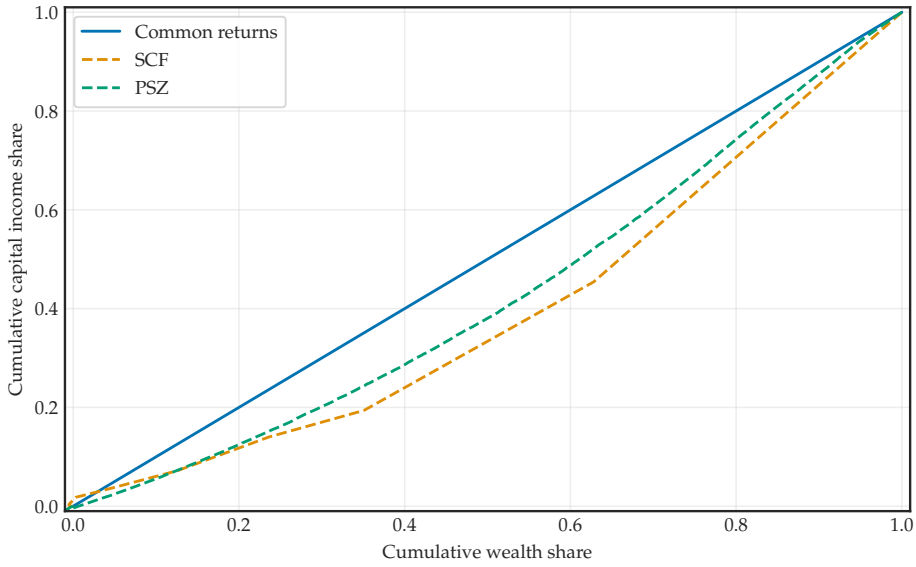


Figure 2: Shares of Wealth and Capital Income

Note: The figure shows the shares of wealth and capital income held by households in the 2019 SCF and the 2019 Distributional National Accounts due to PSZ. Households are sorted by their level of wealth. The x-axis shows the share of wealth held by the bottom p% of households, while the y-axis shows the share of capital income held by the same group.

Figure 2 also makes an additional point: Assuming common returns across households will understate the income of the rich, as they earn higher returns. Thus, assuming common returns is also likely to understate wealth inequality. Gaillard et al. (2023) make a very similar point theoretically, arguing that scale-dependent returns are necessary to match the tops of the distributions of wealth and income. Guvenen et al. (2023) find a similar result. Thus, I conclude that matching the degree of return heterogeneity—and, in particular, that rich households earn higher returns—is crucial for matching the degree of wealth inequality, particularly at the top.

2.2.2 Heterogeneous Returns in the PSID

Having looked at the cross-sectional data, I now take a different approach: Constructing data on heterogeneous returns directly using panel data. In particular, I use the PSID conducted from 1999 to 2019.[6] The panel is biennial, and the unit of

6. The resulting data is for every other year from 2000 to 2018, see Appendix A.4.1.

measurement is households. I detail the data construction in Appendix A.4.[7]

The main outcome is the return on wealth for household i at time t, which is

$$r_{it}^{a} = \frac{x_{it}}{a_{it-1}}, \tag{2}$$

where a_{it-1} is wealth and x_{it} is the income generated from this wealth, both realized and unrealized. Crucially, the return is measured *net of investment*. This is often not possible in other studies because net investment is unknown. For instance, the seminal contribution of Fagereng et al. (2020) uses Norwegian register data, which does not include net investment. While register data has clear advantages, such as high quality, without data on net investment, the measure of return is potentially biased. Consider, for instance, a household that buys a large amount of stocks in between measurements of wealth. The increase in wealth is then falsely counted as capital gains, biasing upward the measure of returns incorrectly. Fagereng et al. (2020) employ approximation methods to minimize the bias from this source. In the PSID, households are asked directly about net investments for many asset categories, so I avoid this potential source of bias.

Capital income, x_{it}, can be split into eight sources: Trust fund and royalties, interest, dividends, primary housing, other housing, businesses, stocks, and other. I discuss how I measure these in Appendix A.4.3. Total wealth, a_{it}, can be split into eight asset categories: Primary housing, other housing, business, stocks, private annuities or IRAs, checking/savings accounts, vehicles, and other assets, cf. Appendix A.4.4. I annualize the returns from the biennial panel. Thus, the resulting measure of returns is *real pre-tax annual returns*. For wealth, I normalize by the average level of wealth in the year. I report descriptive statistics of the data in Appendix A.4.5.

Having constructed a dataset of heterogeneous returns, I start by plotting a histogram of these in Figure 3. As the figure shows, there is significant dispersion in returns. In particular, the standard deviation of returns is 13.3%. This is perfectly consistent with the literature, which finds standard deviations in the range of 7–31 pp. across settings, approaches, and datasets (Bach et al. 2020, Fagereng et al. 2020, Smith et al. 2022, and Snudden 2021). One might ask if the heterogeneity in returns is explained by the types of assets held by households. I show that this is not the case in

7. I am thankful to Stephen Snudden for making his replication files for Snudden (2021) publicly available. My return measurements take a starting point in his work but are different for numerous reasons, so any errors are purely mine.

Appendix A.4.6: Almost all heterogeneity in returns is within asset categories.

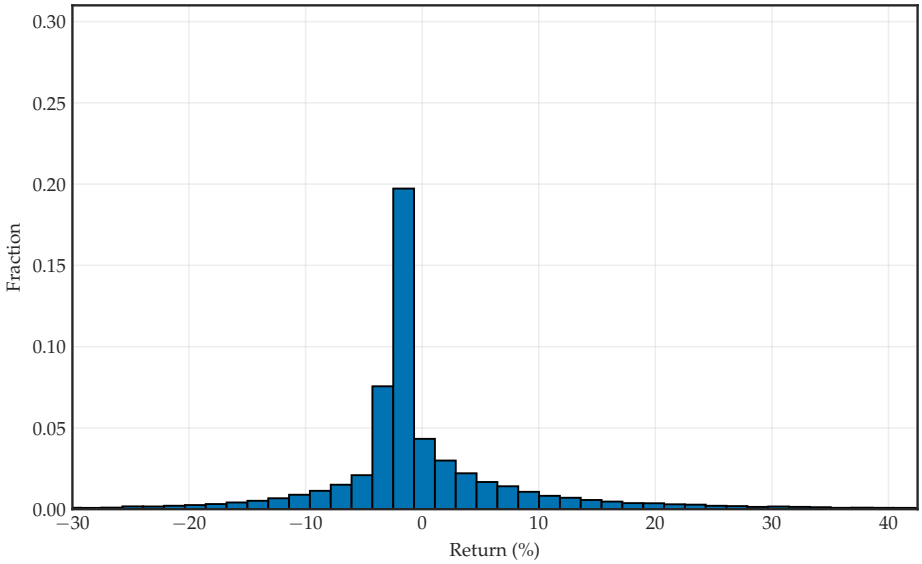


Figure 3: Distribution of Returns

Note: The figure shows a histogram of the returns of US households in the dataset constructed from the PSID.

2.3 Heterogeneous Returns Pass-Through

One thing is the cross-sectional dispersion in returns. Another is how returns change in response to aggregate shocks. To get at this, I estimate the following regression:

$$r_{it}^{a} = \alpha^{(q)} + \beta^{(q)} \bar{r}_{t}^{a} + \varepsilon_{it}^{(q)}, \tag{3}$$

where $\bar{r}_t^a$ is the average returns across households in year t and superscript (q) indicates different deciles of wealth. This is very similar to the idea of measuring the "β" of earnings, i.e., how much households' earnings change as aggregate earnings or GDP changes—see, for instance, Guvenen et al. (2017). However, to the best of my knowledge, I am the first to do this for returns.

Figure 4 plots the estimated β's at different points in the wealth distribution. The figure clearly shows that the pass-through of average to households' returns is stronger for richer households. In particular, the pass-through is nonexistent for the poorest 40% of households, while it is almost 2 for the richest 30%.

How robust are the results in Figure 4? I consider robustness of this in Appendix A.4.7. I find that the results are robust: The returns change more for wealthier households than for poor households when the average return changes. Additionally,

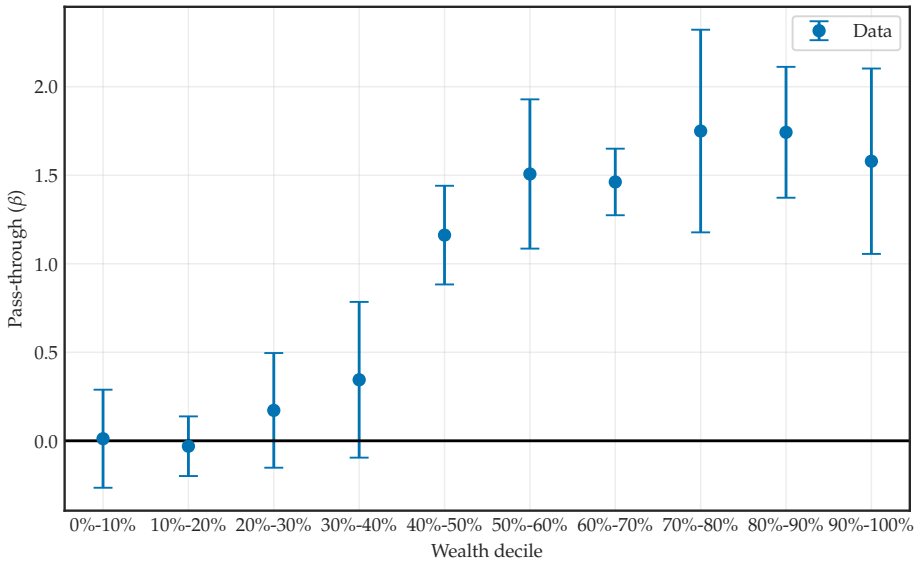


Figure 4: Pass-Through of Average to Households' Returns by Wealth

Note: The figure shows estimates of β from eq. (3) by wealth, a_{it-1}. The standard errors are clustered by year. Confidence intervals are at the 95% level.

I find that the results are robust to clustering the standard errors by household. Finally, I find that results are robust to excluding housing and vehicles and only looking at financial assets, cf. Appendix A.4.8.

In addition to Figure 4, I plot the time series of returns for the bottom 20% and top 20% as well as the average return in Figure 5. The figure clearly shows that the returns for the bottom 20% are mostly flat over the business cycle, while the returns for the rich move more than 1-for-1 with the average return.

What could be driving the different exposure of household levels to the aggregate returns found in Figures 4 and 5? Appendix A.4.9 shows that wealthier households (i) hold riskier assets but (ii) still earn a higher return adjusted for risk.

I explore a different approach to measuring the co-movement of returns in Appendix A.5. Here, I use portfolio shares in the SCF and aggregate returns measured by Jordà et al. (2019). With this different approach in a different dataset, I confirm that the returns for rich households co-move more than 1-for-1 with the average return.

3 A Model of Heterogeneous Returns

Having established that returns are heterogeneous using two different approaches with two different datasets, I now introduce heterogeneous returns in an otherwise

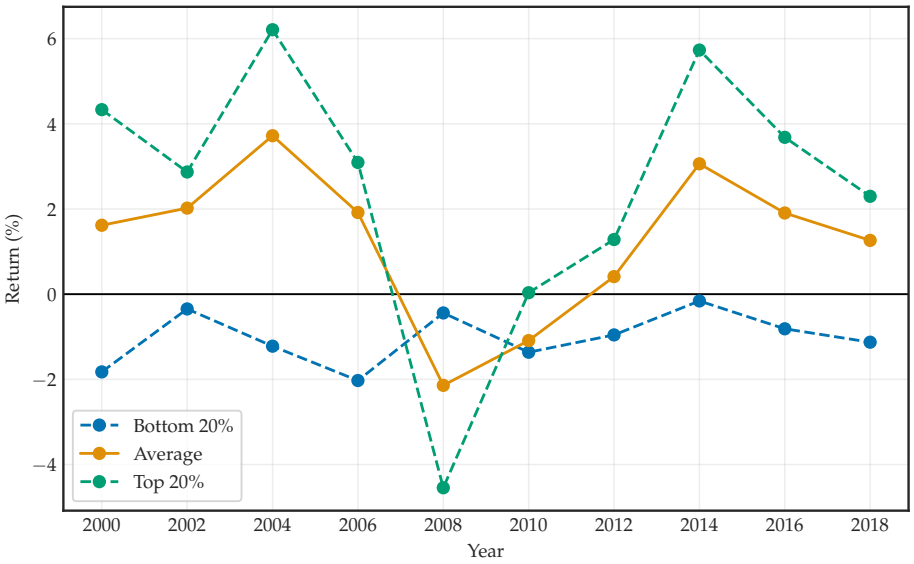


Figure 5: Time Series of Average Returns

Note: The figure shows a time series of average returns based on the PSID data.

standard HANK model. The goal of doing so is both to match the degree of heterogeneity in returns, but also to fit the concentration of wealth at the top.

The model economy consists of households, firms, and a public sector. Households consume and save in an asset that pays back heterogeneous returns. Labor supply is set on behalf of households by a union subject to adjustment costs. Firms use only labor to produce goods under monopolistic competition and flexible prices. The government issues bonds, raises taxes, pays transfers, and consumes. The central bank sets the real interest rate on bonds. The key innovations compared to the literature are how to incorporate heterogeneous returns in the New Keynesian framework. This amounts to changing both the asset demand and supply sides.

I write and solve the model with perfect foresight over aggregates. For small shocks, this is equivalent to solving the model with aggregate risk. I discuss the solution method in more detail in Appendix B.1.

3.1 Households

Time is discrete and the horizon is infinite: $t = 0, 1, \ldots$ There is a continuum of households indexed by $i \in [0, 1]$. Each household chooses the sequence of consumption,

$(c_{it})_{t=0}^{\infty}$, and wealth, $(a_{it})_{t=0}^{\infty}$, to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_{it}) - v(n_{it}) \right],$$

where u and v are the instantaneous utility of consumption and the disutility of labor supply, respectively. $\beta \in [0,1]$ is the common discount factor for all households. I give the recursive formulation of the problem in Appendix B.2. Labor supply is identical for all households and is not chosen directly by the households, $n_{it} = N_t$. Instead, it is chosen by the union as I describe later. Thus, the disutility of labor does not affect household behavior. I consider a standard CRRA utility function,

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma},$$

where $\sigma > 0$ is the inverse elasticity of intertemporal substitution and $u(c) = \log c$ when $\sigma = 1$. Each household is subject to a budget constraint in every period,

$$c_{it} + a_{it} = (1 + r_{it}^a) a_{it-1} + z_{it} + T_t - t_{it}, \tag{4}$$

where r_{it}^a are the heterogeneous returns, z_{it} is the real pre-tax labor earnings, T_t is lump-sum transfers, and t_{it} is the tax bill. With this budget constraint, let me already discuss what determines the distributional effects of aggregate shocks to motivate the rest of the model. In particular, I derive sufficient statistics for the distributional effects of shocks, in the sense that any model that matches these will yield the same distribution of income in response to the shock. The components of the sufficient statistics are given in Proposition 1.

Proposition 1. *Consider an aggregate shock affecting returns and labor income. Let $dX_t \approx X_t - X_{ss}$ denote the perturbation around steady state of any variable X. The share of income generated by this shock going to household i on impact is*

$$\frac{d\psi_i}{d\Psi} = \alpha_z \frac{z_i}{Z} \frac{d \log z_i}{d \log Z} + (1 - \alpha_z) \frac{a_{i,-1}}{A_{-1}} \frac{dr_i^a}{d\tilde{r}^a}, \tag{5}$$

where $X = \int x_i \, di$ for any x, $\alpha_z = dZ/d\Psi$ is the labor share of monetary policy, $\psi_{it} = r_{it}^a a_{it-1} + z_{it} + T_t$ is income, and $\tilde{r}^a = \int r_i^a a_{i,-1} \, di / A_{-1}$ and $Z = \int z_i \, di$ are aggregate returns and earnings.

Proof. See Appendix B.3. □

This shows that one aggregate component and four micro components determine the distribution of income generated by shocks such as monetary policy. The aggregate component is the capital income share, $1 - \alpha_z$, which motivates focusing on this quantity. The four micro components in the sufficient statistics decomposition are: The wealth distribution, $a_{i,-1}$, the labor income distribution, z_i, and the pass-through of aggregates to these. For this reason, I emphasize these components both in the construction and calibration of the model. Let me start by discussing how I specify the earnings and returns processes for this purpose.

Real pre-tax labor earnings, z_{it}, depends on an idiosyncratic component and an aggregate component,

$$z_{it} = e^z_{it} Z_{ss} + e^z_{it} \beta^z_{it} (Z_t - Z_{ss}),$$

where $\int e^z_{it}\, di = 1$ and $\int e^z_{it} \beta^z_{it}\, di = 1$ such that Z_t is the average earnings. The idiosyncratic component of earnings follows a Markov chain,

$$e^z_{it} \sim \text{Markov}(\mathcal{S}_z, \mathcal{P}_z),$$

where $\mathcal{S}_z$ is the state space and $\mathcal{P}_z$ is the transition matrix. β^z_{it} measures the elasticity of households' earnings to aggregate earnings, i.e., the "worker β" in the style of Guvenen et al. (2017), cf. Appendix B.4. It is calibrated to empirical evidence in Section 4. $\beta^z_{it} = 1$ nests the standard specification of $z_{it} = e^z_{it} Z_t$.[8]

The return on wealth, r^a_{it}, has idiosyncratic and aggregate components,

$$r^a_{it} = r^a_{ss} + e^r_{it} + \beta^r_{it} (r^a_t - r^a_{ss}), \tag{6}$$

where $\int e^r_{it}\, di = 0$ and $\int \beta^r_{it}\, di = 1$ such that r^a_t is the average return. The idiosyncratic component follows a Markov chain,

$$e^r_{it} \sim \text{Markov}(\mathcal{S}_r, \mathcal{P}_r),$$

where $\mathcal{S}_r$ is the state space and $\mathcal{P}_r$ is the transition matrix. Thus, the return on wealth is heterogeneous across households due to the randomness in e^r_{it}. Furthermore, returns may be persistent: If households earn a high return in one period, they tend to also earn a high return in the next period. β^r_{it} controls the pass-through of aggregate to

8. In this case, the pass-through is common: $\partial \log z_{it} / \partial \log Z_t = 1$.

households' returns. I specify the functional form of this in the calibration. A relevant special case is equal pass-through, i.e., $\beta^r_{it} = 1$ for all i and t.

Households are taxed on both capital and labor income at rate τ_t:

$$t_{it} = \tau_t \left(r^a_{it} a_{it-1} + z_{it} \right).$$

Finally, households are subject to a borrowing constraint:

$$a_{it} \geq 0.$$

In this paper, I will focus on the case with heterogeneous returns, i.e., $e^r_{it} \neq 0$. Let me briefly mention a special case that I will compare the model to. This is the standard HANK model. This is nested when returns are common, $e^r_{it} = 0$ and $\beta^r_{it} = 1$.

3.2 The Rest of the Model

Having presented the household side with heterogeneous returns, let me now present the rest of the model. For the most part, this is the standard New Keynesian model. The key difference is where the heterogeneous returns come from.

3.2.1 Firms

Firms produce output Y_t using only labor N_t with constant returns to scale, $Y_t = N_t$. I consider capital in Section 6.3. They sell this output to households at price P_t and pay households a wage rate W_t for their labor, such that real labor income is $Z_t = w_t N_t$, where $w_t = W_t / P_t$ is the real wage rate. Firms set prices flexibly, P_t, at a markup, $\mu \geq 1$, over marginal costs:

$$P_t = \mu W_t. \tag{7}$$

Instead of sticky prices, the nominal rigidity is in the form of sticky wages, as is standard in the HANK literature. I consider sticky prices in Section 6.3. All profits are paid period-by-period to households as dividends, which in real terms are

$$D_t = Y_t - Z_t.$$

Firms issue a unit mass of shares, which they sell at real price p_t.

3.2.2 Government

The government issues real bonds, B_t, which pay real interest rate r_t. It uses the bonds and taxes, $\mathcal{T}_t$, to finance government consumption, G_t, and lump-sum transfers to households, T_t, such that the real budget constraint is

$$B_t = (1 + r_{t-1})B_{t-1} + G_t + T_t - \mathcal{T}_t, \tag{8}$$

where tax receipts are $\mathcal{T}_t = \int t_{it}\, di$.

Shocks to government consumption or transfers are financed by both issuance of bonds and higher taxes in the short run. In the long run, fiscal policy is passive, in the sense that the tax rate, τ_t, adjusts to ensure that

$$B_t = B_{ss} + \phi_B(B_{t-1} - B_{ss}) + (G_t - G_{ss}) + (T_t - T_{ss}) \tag{9}$$

following Auclert et al. (2024a). The rule is chosen to ensure that bonds return to the initial steady state in the long run, i.e., $\lim_{t\to\infty} B_t = B_{ss}$, where "ss" indicates the steady state of the model.

3.2.3 Central Bank

The central bank sets the real interest rate directly,

$$r_t = r_{ss} + \varepsilon_t, \tag{10}$$

where ε_t is a monetary policy shock. The nominal interest rate is then

$$i_t = (1 + r_t)(1 + \pi_{t+1}) - 1,$$

where $\pi_t = P_t/P_{t-1} - 1$ is inflation.

3.2.4 Union

In addition to sticky prices, there are also nominal rigidities in the model in the form of sticky wages. Specifically, a union sets nominal wages subject to Rotemberg adjustment costs. This yields the following non-linear New Keynesian wage Phillips

curve (NKWPC) as in Auclert et al. (2024b):

$$\pi_t^W(1+\pi_t^W) = \kappa^W\left(\frac{v'(N_t)}{u'(C_t)(1-\tau_t)w_t} - 1\right) + \frac{1}{1+r_t}\pi_{t+1}^W(1+\pi_{t+1}^W). \tag{11}$$

This NKWPC is written in such a way that heterogeneity does not matter directly as is common in the HANK literature.[9] $v(N_t) = \gamma N_t^{1+1/\phi}$ is the disutility of labor with $\phi > 0$ measuring the Frisch elasticity of labor supply and γ being a scalar parameter.

3.2.5 Asset Supply

Households do not have a portfolio choice but instead choose overall savings, a_{it}, with return r_{it}^a. The savings reflect two assets in the economy: Firm equity and government bonds. Both assets pay the same return along the perfect foresight transition path:[10]

$$\frac{p_{t+1} + D_{t+1}}{p_t} - 1 = r_t. \tag{12}$$

These returns are paid to households according to the heterogeneous returns process in the household problem. This means that the total capital income coming from heterogeneous returns for all households equals the capital income in the economy:[11]

$$\underbrace{\int r_{it}^a a_{it-1}\, di}_{\text{HH's capital income}} = \underbrace{p_t + D_t + (1+r_{t-1})B_{t-1} - A_{t-1}}_{\text{Capital income from bonds and equity}}. \tag{13}$$

This defines the average return in the economy. Intuitively, the average return adjusts each period to ensure that capital income in general equilibrium matches the capital income earned by households, taking into account the idiosyncratic returns.

It becomes convenient to define the *wealth-weighted* average return,

$$\tilde{r}_t^a = \frac{\int r_{it}^a a_{it-1}\, di}{A_{t-1}},$$

where $A_t = \int a_{it}\, di$. Note that the wealth-weighted average return generally differs

9. This simplifies the comparison of different HANK models as only the household side is affected while the NKWPC is unchanged.

10. The exception is on impact, $t = 0$, where the shock causes an unexpected revaluation of assets.

11. By capital income, I mean any income generated by holding assets. This could also be called asset income. I consider physical capital in Section 6.3.

from the average return since

$$\tilde{r}_t^a = r_t^a + \underbrace{\text{Cov}\left(r_{it}^a, \frac{a_{it-1}}{A_{t-1}}\right)}_{\text{Scale dependence}}.$$

The covariance measures the degree of scale dependence in returns: The degree to which rich households earn high returns on their wealth.[12] Intuitively, *if* rich households earn higher returns, i.e., $\text{Cov}\left(r_{it}^a, \frac{a_{it-1}}{A_{t-1}}\right) > 0$, the wealth-weighted average return is higher. In the case of common returns, the wealth-weighted average return and average return are identical, $\tilde{r}_t^a = r_t^a$, since $\text{Cov}\left(r_{it}^a, \frac{a_{it-1}}{A_{t-1}}\right) = 0$ as $r_{it}^a = r_t^a$.

3.2.6 Market Clearing

Final output is consumed by households and the government,

$$Y_t = C_t + G_t, \tag{14}$$

where private consumption is aggregated over households, $C_t = \int c_{it}\, di$. Asset market clearing, $A_t = \int a_{it}\, di = p_t + B_t$, follows by Walras' law, cf. Appendix B.6.

3.2.7 Solution and Equilibrium

Introducing heterogeneous returns makes solving the model non-standard. I discuss how I solve the model in Appendix B.1, where I also define the equilibrium concept.

4 Calibration

I now present the calibration of the model. I first discuss the calibration of the new part, the returns process, to the panel of heterogeneous returns in Section 2. I then proceed with the calibration of the rest of the model, which is quite standard.

12. Others use "scale dependence" to mean that wealthier households earn higher returns *because* they are wealthier. I use the term scale dependence more generally to refer to a positive covariance between wealth and returns. This might occur directly *because* the returns process is increasing in wealth—as I consider in Appendix B.5—or it might occur endogenously, as in my baseline.

4.1 The Returns Process

I start by calibrating the returns process to my data on household-level returns in the US. There are three key objects to be calibrated: The average return, r^a_{ss}, the state space, $\mathcal{S}_r$, and the transition matrix, $\mathcal{P}_r$. Let me go through each.

I simply calibrate the average return, r^a_{ss}, to the average return in the data.[13] I then calibrate the idiosyncratic component of returns, e^r_{it}. This process is controlled by the state space, $\mathcal{S}_r$, and the transition matrix, $\mathcal{P}_r$. I start by constructing the equivalent of the idiosyncratic component of returns in the data. In particular, I care about the component of returns that is not permanent to the household and that is not driven by aggregate movements in returns. For this reason, I estimate a regression of returns on a household fixed effect, time fixed effect, and age dummies.[14] I use the residual from this regression as my measure of idiosyncratic returns. I divide the data on residual returns into 7 bins and compute the median return within each bin. This yields the stationary distribution shown in Figure 6.

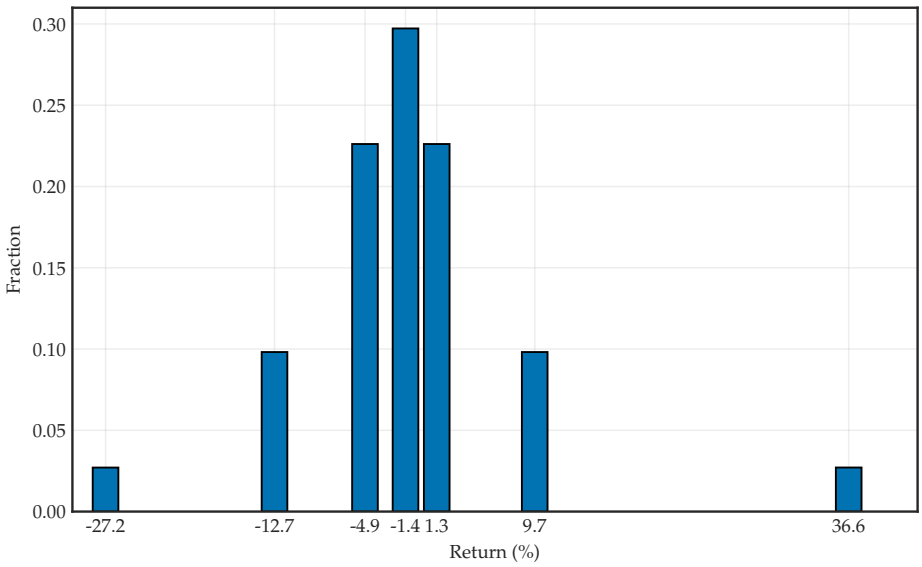


Figure 6: Histogram of the Returns in the Model

Note: The figure shows a histogram of the cross-sectional distribution of heterogeneous returns, r^a_{it}.

Next is the transition matrix. I assume that households either stay in their current

13. The empirical return conditions on $a_{it-1} > 0$. Otherwise, the denominator when computing returns is 0. Thus, I set r^a_{ss} so the average return in the model, conditional on $a_{it-1} > 0$, matches the data.

14. I estimate $r^a_{it} = \gamma_i + \gamma_t + \beta D_{it} + \varepsilon_{it}$, where D_{it} is a vector of age dummies and ε_{it} is the residual.

state or increase or drop one state. Additionally, the probability of changing state is the same whether going up or down. After matching the stationary distribution, this leaves one degree of freedom. I use this to match the number of billionaires in the US. The resulting annual autocorrelation is 0.69. I discuss the returns process in more detail in Appendix C.1. I verify that the returns process is such that the model has a stationary distribution in Appendix C.2.

Finally, I specify a functional form of β^r_{it} to match the pass-through of average to households' returns in the data.[15] In particular, I set

$$\beta^r_{it} = \frac{\log\left\{1 + (1+\theta_0)(1-\theta_1^{a_{it-1}})^2\right\}}{\log(2+\theta_0)} \quad \text{and} \quad \hat{\theta} = \arg\min_{\theta} \frac{1}{G}\sum_{g=1}^{G} \omega_g \left(\beta^g_{\text{data}} - \beta^g_{\text{model}}\right)^2$$

where ω_g are weights equal to the inverse of the standard errors. I estimate $\hat{\theta} = (230, 0.19)'$, with the resulting fit shown in Figure 7.[16] I consider robustness to the functional form in Appendix C.3. I find that any functional form that yields a similar fit yields almost numerically identical results from the chosen functional form.

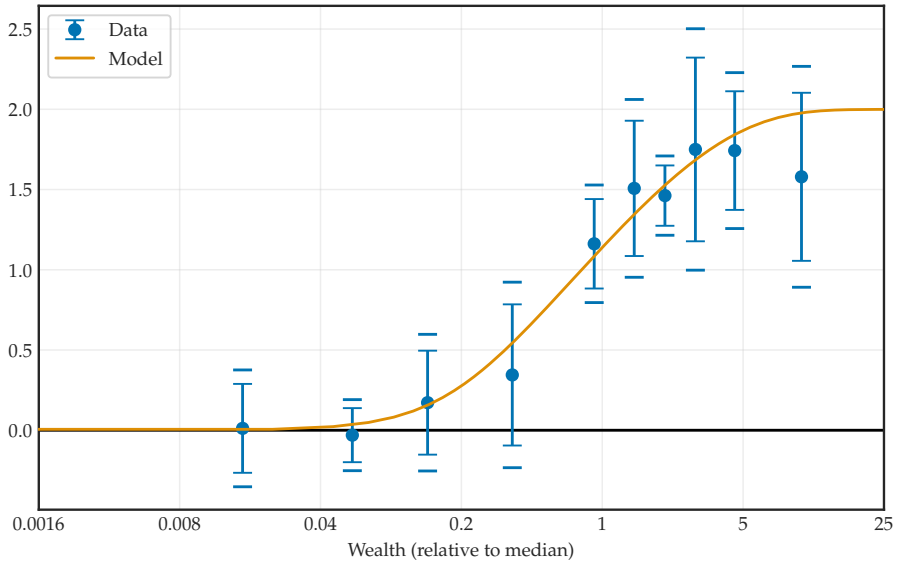


Figure 7: Pass-Through of Aggregate to Households' Returns, β^r_{it}

Note: The x-axis shows $a_{i,-1}$ (relative to the median in the model), while the y-axis shows β^r_{it}. For the data, each point is the median within 10 sorted buckets of 10% of the wealth distribution. To allow comparing the model and the data, the data has been scaled such that the median wealth is the same as in the model. Ticks indicate 95% and 99% confidence intervals.

15. I normalize this so $\int \beta^r_{i,ss}\, di = 1$. I also add a tiny value to a_{it-1} in the exponent to avoid 0^0.

16. I do *not* condition on $a_{it-1} > 0$ in Figure 7. If I did, the model-implied β^r_{it} would tend to be higher than in the data, because $\int \beta^r_{it}\, di = 1$ is computed over *all* households, not just the ones with $a_{it-1} > 0$.

4.2 The Rest of the Model

I calibrate the model to match the US in 2019 at an annual frequency. For the household side, I mainly use the 2019 edition of the SCF.[17] I start by presenting the internally calibrated parameters outside the returns process. These are parameters set to match moments in the data. I calibrate the discount factor to match the asset demand-to-GDP ratio of $\int a_{i,ss}\, di/Y = 447\%$, which yields $\beta = 0.950$.

	Description	Value	Target	Source
β	Discount factor	0.950	$\int a_{i,ss}\, di/Y = 447\%$	SCF 2019
r^a	Average return	-1.3%	$\mathbb{E}[r^a_{it} \mid a_{it-1} > 0]$	PSID
T	Transfers	0.17	$T/\text{Income} = 14\%$	SCF 2019
B	Government bonds	1.05	$B/Y = 105\%$	US OMB
τ	Tax rate	0.38	$G/Y = 17.6\%$	US BEA
μ	Markup	1.19	$A/Y = 447\%$	SCF 2019
c	Return mobility parameter	0.024	788 billionaires	Henley
σ_z	Log earnings std. dev.	0.561	Bot. 80% earnings $= 35\%$	SCF 2019
ρ_z	Log earnings persistence	0.91	Floden and Lindé (2001)	
σ	CRRA	1	Kaplan et al. (2018)	
ϕ_B	Tax adjustment speed	0.9	Auclert et al. (2024a)	
κ^W	NKWPC slope	0.03	Auclert et al. (2024b)	

Table 1: Calibration

Note: "Income" refers to $\Psi_t = Z_t + \int r^a_{it} a_{it-1}\, di + T_t$.

I let the earnings process be an AR(1) in logs, which is discretized using the Rouwenhorst method (Kopecky and Suen 2010 and Rouwenhorst 1995). I calibrate the autocorrelation of log earnings to $\rho_z = 0.91$ as in Floden and Lindé (2001). I calibrate

17. The data follows Kuhn and Rios-Rull (2016) updated to the 2019 SCF.

the standard deviation of log earnings to match the earnings share of the bottom 80%. This yields $\sigma_z = 0.561$, which is slightly higher than Guvenen et al. (2021).

For transfers from the government to households, I match an average transfer income share of 14% from the SCF, yielding $T = 0.17$. For government bonds, I match the debt-to-GDP ratio of 105% in 2019 by setting $B = 1.05$. I match the ratio of government consumption to GDP of 17.6%, which requires setting $\tau = 0.38$, essentially matching the estimates in Barro and Redlick (2011). Finally, I calibrate the markup such that the supply of assets matches the demand for assets, yielding $\mu = 1.19$.

Next, I present the externally calibrated parameters, i.e., the parameters set to values from the literature. Here, I set a CRRA of $\sigma = 1$, i.e., I consider log utility, which is standard. I set the tax adjustment speed to $\phi_B = 0.9$ as in Auclert et al. (2024a). I set the slope of the Phillips curve, κ^W, in accordance with Auclert et al. (2024b).

Finally, let me discuss the calibration of β^z_{it}: The pass-through of aggregate earnings to households' earnings. For this purpose, I calibrate to the "worker β's" from Guvenen et al. (2017). In particular, I set β^z_{it} as a function of z_{it} cf. Appendix B.4. The result is given in Figure 8. The figure shows a U-shaped relationship between earnings and the pass-through of earnings: The pass-through is largest for the rich and the poor, while the pass-through is low for the middle of the distribution.

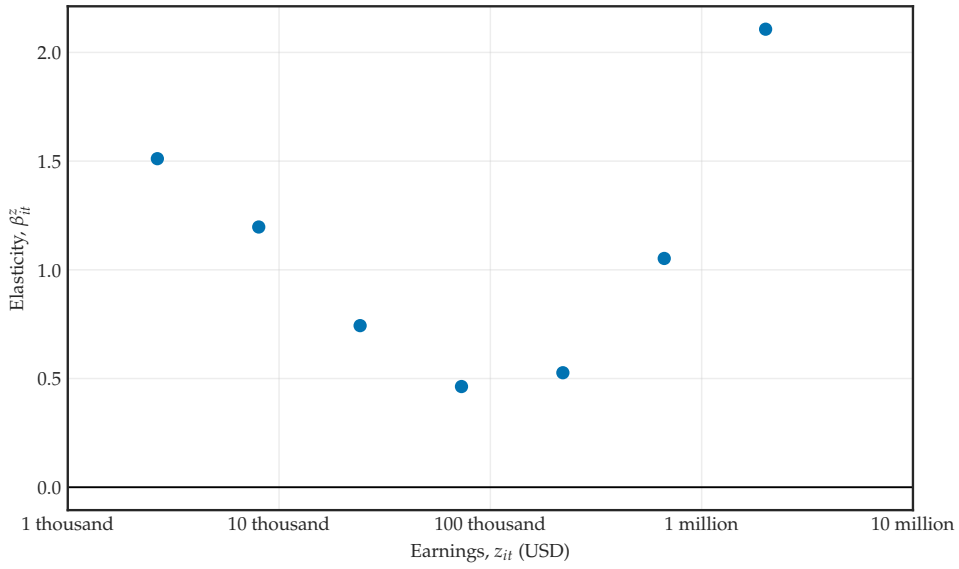


Figure 8: Pass-Through of Average to Households' Earnings by Earnings, β^z_{it}

Note: The figure shows the log pass-through of Z_t to z_{it} by values of z_{it}.

In the rest of the paper, I compare the model with heterogeneous returns to the

standard HANK model. Table 1 is for the model with heterogeneous returns, so let me discuss the calibration of the standard HANK model. All parameters are the same except the following. I set $e^r_{it} = 0$ and $\beta^r_{it} = 1$ to get common returns. Additionally, I consider permanent discount factor heterogeneity: Half of all households have a discount factor $\overline{\beta}$, while the other half have a discount factor $\underline{\beta}$. The first discount factor is set to match asset demand—as in the model with heterogeneous returns. The other discount factor is set to match the same average marginal propensity to consume (MPC) as the model with heterogeneous returns. As I argue shortly, this is an important moment to match, which the standard HANK model cannot do without permanent discount factor heterogeneity (or some other change).

5 Microeconomic Fit

In this section, I show how the model with heterogeneous returns replicates several aspects of the microeconomic data that standard HANK models do not. To do so, I start by discussing heterogeneous returns and then turn to the wealth distribution.

5.1 Heterogeneous Returns

I now consider the heterogeneous returns in the model and their fit to the data. As discussed in Section 4, the returns process is calibrated to match the cross-sectional dispersion of returns. Thus, the model, by assumption, matches this aspect.

A more interesting feature of the returns in the model is that they exhibit *scale dependence*. Scale dependence is the observation that rich households tend to earn higher returns. This is well-documented in the literature, see Fagereng et al. (2020), Bach et al. (2020), Xavier (2021), and Daminato and Pistaferri (2024). To show this, I split up the data sample and the model distribution by the wealth levels of households. For each group, I then compute the average return. Figure 9 plots this for quintiles of wealth. The figure shows that households with more wealth tend to earn higher returns, both in the model and in the data. The relationship between returns and wealth is slightly stronger in the model than in the data. In this regard, the model is more consistent with the literature as well as my data when only looking at financial returns, which both feature stronger relationships, cf. Appendix D.1.

Note that there is nothing in the specification of the returns process that says that rich households earn higher returns. Instead, the relationship between returns and

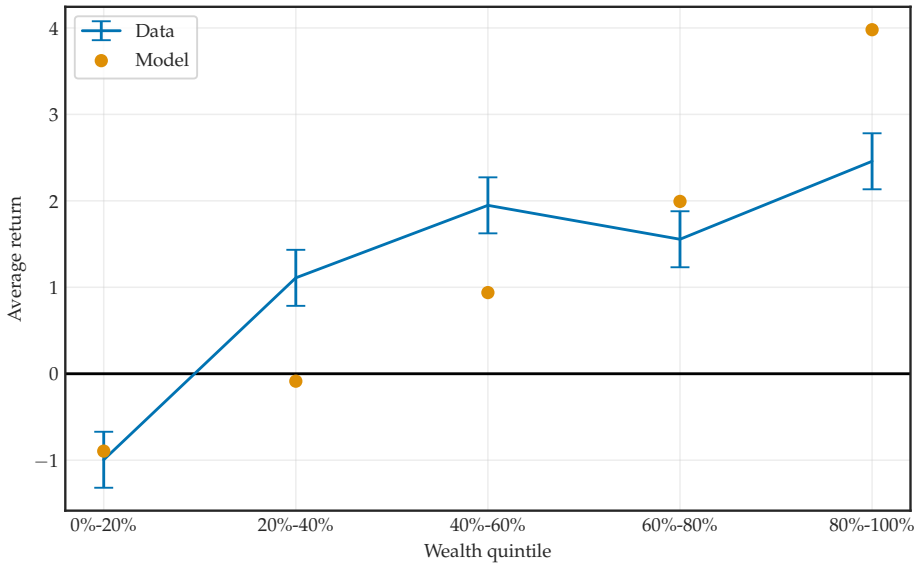


Figure 9: Average Return by Wealth (Scale Dependence)

Note: The figure shows the average return by quintiles of a_{it-1}.

wealth comes about as a result of household decisions: If households earn higher returns, they choose to save more, creating the relationship in Figure 9. I study this further in Appendix D.2, which plots the policy consumption functions for different rates of returns. The figure shows that households that earn higher returns have a lower MPC, saving more of income windfalls instead.

To study the congruence of return heterogeneity in the model and the data further, Figure 10 shows the equivalent of Figure 2 in the models. The figure shows that the bottom p% of households in the wealth distribution always hold the same shares of wealth and capital income in the standard HANK model, i.e., the curve is on the 45-degree line. This is due to the model having common returns. This clearly is at odds with the data, where, for instance, the bottom 95% holds 35% of all wealth but only around 20% of capital income. In contrast, the model with heterogeneous returns fits the data well due to heterogeneity in returns and the scale dependence of returns.

5.2 Wealth

I now turn my attention to the wealth distribution and, in particular, the concentration in the upper right tail. I plot the right tail of the wealth distribution in Figure 11. Importantly, if wealth is Pareto-distributed, this plot should look like a line according to eq. (1), as I discussed in Section 2.1. As mentioned in Section 2.1, the two HANK models in Kaplan et al. (2018) and Auclert et al. (2024b) provide a poor fit to the

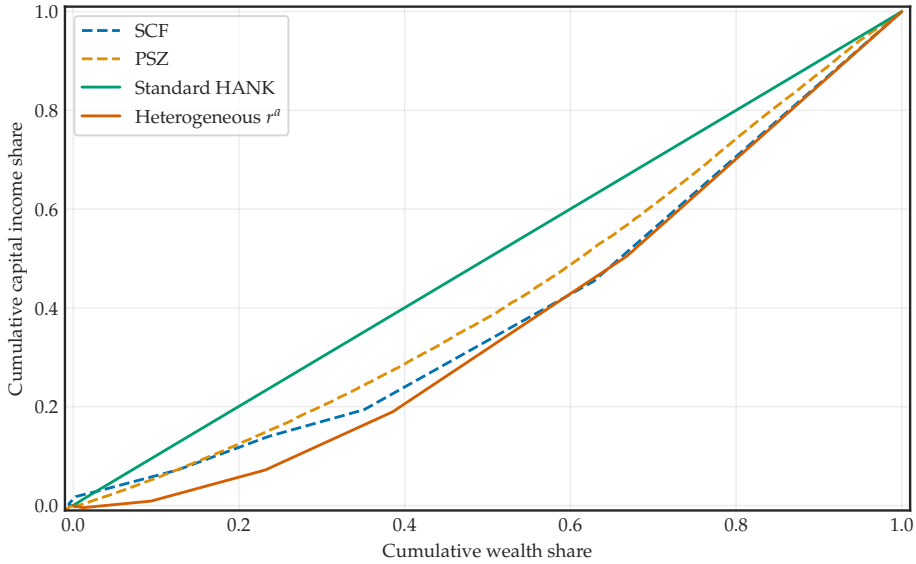


Figure 10: Shares of Wealth and Capital Income in the Models

Note: The figure shows the shares of wealth and capital income at different points in the distribution of households in the 2019 SCF, the 2019 Distributional National Accounts due to PSZ, and the two models. Households are sorted by their level of wealth. The x-axis shows the share of wealth held by the bottom p% of households, while the y-axis shows the share of capital income held by the same group.

top of the wealth distribution. In addition to these models, Figure 11 also plots the "standard HANK" model, i.e., the model with common returns. However, this model also understates top wealth. On the other hand, the model with heterogeneous returns replicates the data almost perfectly, displaying a Pareto tail.

To elaborate on Figure 11, I provide key statistics on the wealth distribution in Table 2. This table highlights the key conclusion: The inclusion of heterogeneous returns allows the model to replicate the wealth distribution. Particularly, the match to the top of the wealth distribution is a key innovation, as this is known to be difficult in standard heterogeneous agent models. This is even though only the number of billionaires in Table 2 is targeted in the calibration of the model. The improvement of the fit to the wealth distribution is not at the cost of matching the share of hand-to-mouth households, which is around 1/3 in both models. In Appendix D.3, I show the Lorenz curve for wealth along income and earnings. This shows that the fit to the whole wealth distribution is good. Additionally, Appendix D.4 shows that the relative ranking of the concentration of wealth, consumption, and capital income is consistent with the ranking in data, as emphasized by Gaillard et al. (2023).

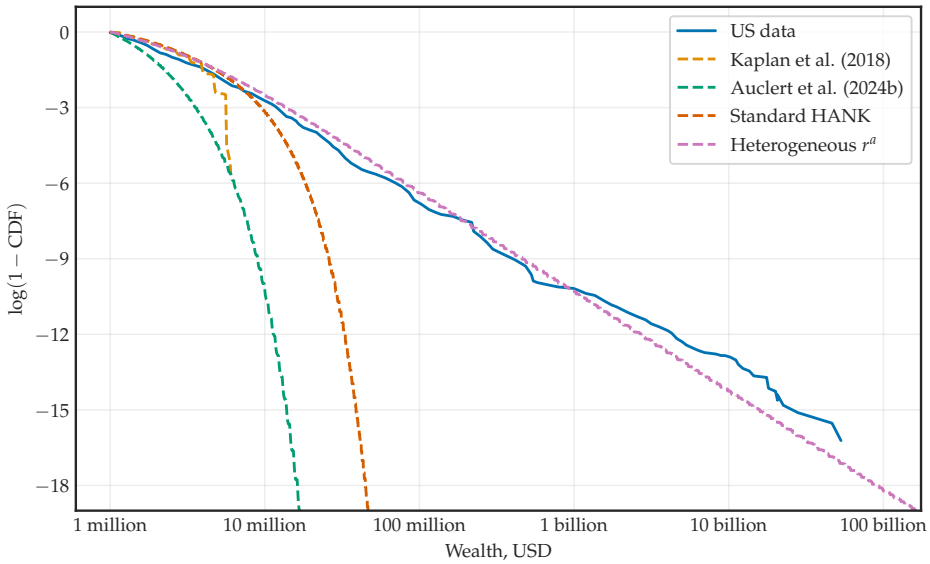


Figure 11: Wealth Concentration at the Top

Note: The figure shows the distribution of wealth in several models. It does so by plotting the counter-CDF against the level of wealth in USD. Both axes are log-scale. The models are the model from Kaplan et al. (2018), the "HA-hi-liq" model from Auclert et al. (2024b), the standard HANK model with common returns, and the model with heterogeneous returns. The figure also shows the counter-CDF for the US based on data from Guvenen et al. (2023) and Vermeulen (2018).

Moment	Standard HANK	Heterogeneous r^a	Data
Top 20% share	85%	90%	87%
Top 10% share	63%	77%	76%
Top 1% share	14%	33%	37%
Top 0.1% share	2%	13%	14%
No. of billionaires	0	788	788
Top 0.0006% cutoff (mil. USD)	30	986	1000

Table 2: The Distribution of Wealth

Note: The data on the wealth distribution is from the 2019 SCF. The data on the number of billionaires is from https://www.henleyglobal.com/publications/usa-wealth-report-2024. The values in the model are converted to USD by multiplying by GDP per household for the US in 2023.

The fact that the model matches the wealth distribution is an attractive property in and of itself. But it is also an attractive property for another reason: It lets the model match a realistically high average MPC *at the same* time it matches a realistically

high average wealth. This is important because the literature on HANK models emphasizes that a key object to match is the MPC (Auclert et al. 2024b). However, it is a well-known issue that standard HANK models struggle to match simultaneously the MPC and a realistic level of wealth.[18] This is the "MPC-wealth" trade-off of Kaplan and Violante (2022). The intuition for this trade-off is straightforward: A high MPC is achieved by having a realistic number of households close to or at the borrowing constraint. In contrast, a high level of wealth is achieved by having households away from the borrowing constraint.

The model with heterogeneous returns significantly reduces this tension. To see this, consider Figure 12. Figure 12 recalibrates the model for different levels of variation in returns.[19] As the figure shows, a larger variation in returns is associated with a higher MPC, keeping the aggregate level of wealth fixed.

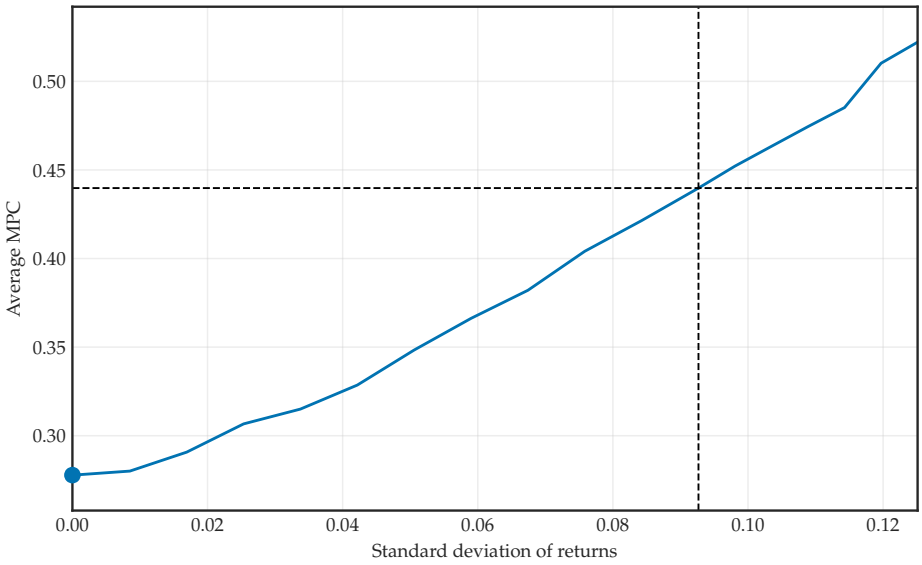


Figure 12: The MPC as Returns Are More Heterogeneous

Note: The figure shows the average MPC for the model with heterogeneous returns as the standard deviation of returns changes. The dashed lines indicate the standard deviation of returns and average MPC in the baseline calibration.

The intuition follows from the considerations regarding the wealth distribution.

18. A realistic annual MPC is around 0.35–0.55. For instance, Fagereng et al. (2021) estimate an MPC of 0.51, while Jappelli and Pistaferri (2014) estimate 0.48. However, Orchard et al. (2025) argue that a lower MPC is realistic. A standard calibration in HANK models is a quarterly MPC of 0.15–0.25, which corresponds to around 0.35–0.55 annually.

19. In particular, I multiply the return grid, $\mathcal{S}_r$, by some varying factor, and keep all other parameters fixed, except β and r, which are recalibrated to ensure that the asset market still clears and the capital income of households still adds up to total capital income.

With a higher variation in returns, the wealth distribution is more spread out. Thus, it is possible to simultaneously match a high MPC as many households are at the borrowing constraint, while matching a high level of wealth, as some ultra-rich households are contributing to a large aggregate wealth. As such, the model with heterogeneous returns is an alternative—and arguably simpler—method of matching the level of wealth and the MPC compared to, for instance, a two-asset model.

6 Implications

In this section, I compare the effects of macroeconomic policies when returns are heterogeneous with a focus on monetary policy. I first focus on the *aggregate* effects, i.e., the effects on aggregate outcomes. Next, I focus on the *distributional* effects, i.e., how different households are affected differently.

6.1 Aggregate Effects of Macroeconomic Shocks

I start by considering the aggregate effects of monetary policy when returns are heterogeneous. To be specific, I consider the economy in the ergodic steady state. I then consider an unexpected 1 percentage point shock to the real interest rate with persistence 0.43, consistent with the estimated monetary policy shock in Auclert et al. (2020). I then report the resulting transition path back to the steady state. I will compare two models: The model with heterogeneous returns and the standard HANK model with common returns. The impulse response functions (IRFs) are given in Figure 13. IRFs for additional variables are given in Appendix E.1.

Figure 13 shows that the effects of monetary policy are very similar in the model with heterogeneous returns compared to the model with common returns. Thus, my model extends the (approximate) neutrality of heterogeneity in the aggregate transmission of monetary policy established in Werning (2015) and Kaplan et al. (2018) to a model with more elaborate heterogeneity. Why is this? This is because whether households are ultra-rich or just "rich", they behave almost according to the permanent income hypothesis. Thus, changing the wealth distribution such that wealth is more concentrated among the ultra-rich instead of the less rich—as the model with heterogeneous returns does—does not drastically change the transmission of monetary policy in the sense that the effects on consumption are very similar. This is consistent with Bilbiie et al. (2025), who argue that heterogeneity does not change the aggregate transmission of shocks much.

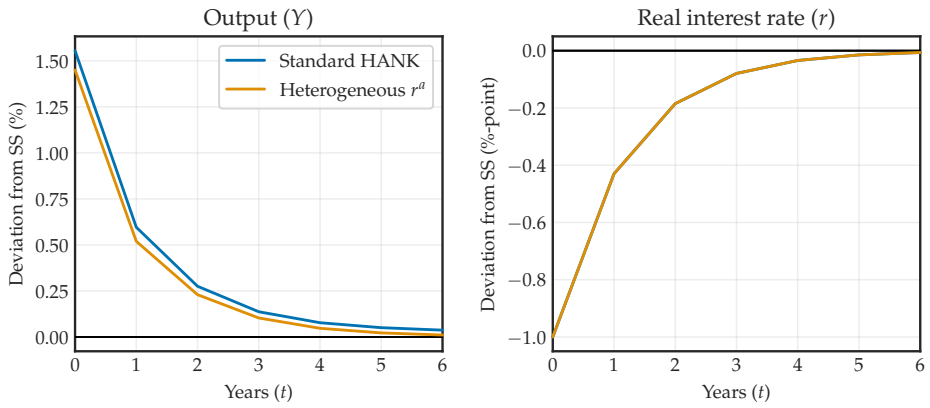


Figure 13: The Aggregate Effects of Monetary Policy

Note: The figure shows impulse response functions to a 1 percentage point fall in the real interest rate. The x-axis shows years after the shock.

While the aggregate effects of monetary policy are similar with and without heterogeneous returns, there is a small difference. This is because heterogeneous returns introduce a new redistribution channel of monetary policy that attenuates its effectiveness. However, this channel is numerically weak, cf. Appendix E.2.

Having understood this, I now turn to capital income. In particular, why does capital income increase in response to a cut in interest rates? The important thing to understand is that monetary policy inflates asset prices. It does so for two reasons: Lower discount rates and higher dividends. To see this, note that eq. (12) implies that the price of firm equity is the discounted sum of future dividends:

$$p_t = \frac{D_{t+1}}{1+r_t} + \frac{D_{t+2}}{(1+r_t)(1+r_{t+1})} + \dots$$

The direct effect of monetary policy is clear: Monetary policy lowers the interest rate, so it decreases the discount rate. This increases the value of any (positive) stream of dividends. This is the main effect, but there is also an indirect effect through dividends: Higher economic activity implies higher dividends.[20] This follows from the fact that wages are more rigid than prices and is consistent with empirical evidence. This makes the indirect effect of monetary policy on asset prices clear: Easy monetary policy increases output and hence dividends, boosting asset prices, cf. Appendix E.1.

In Appendix E.3, I compare my models to the benchmark result from Werning (2015), who shows that monetary policy has the same effects on consumption with

20. This can be seen clearly from the fact that dividends are $D_t = \frac{\mu-1}{\mu} Y_t$.

heterogeneous agents as in a representative agent model under certain assumptions. While this result *approximately* holds in my model, it actually does not hold *exactly*—with or without heterogeneous returns. However, when making certain assumptions, it does hold when comparing the model with common returns to a representative agent model. But—even under these assumptions—the Werning (2015) result still does not hold in the model with heterogeneous returns, which has slightly different effects of monetary policy compared to a model with common returns and a representative agent model, cf. Appendix E.3.

Having considered the effects of monetary policy, I now briefly turn my attention to fiscal policy. In particular, I consider two fiscal policies: Government consumption, G_t, and transfers, T_t. For both these shocks, I consider paths that yield an identical effect on output as with monetary policy in the model with heterogeneous returns.[21] Appendix E.4 discusses how I do this. Doing so is particularly convenient when considering the distributional effects in Section 6.2, as I am holding the aggregate effects across shocks fixed.

Figure 14 shows the output responses in the models with heterogeneous returns and the standard model. The figure shows that the effects of both fiscal policies on output are very similar. This is because both models are calibrated to the same MPC, which is a key determinant of the effectiveness of fiscal policy.

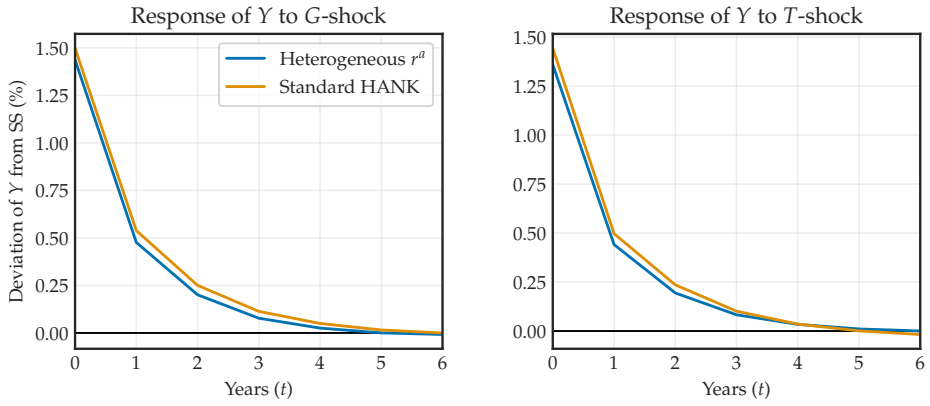


Figure 14: The Aggregate Effects of Fiscal Policy

Note: The figure shows impulse response functions of output to fiscal policy shocks.

21. I use these shocks in both models to compare the effects of the *same* shock on output.

6.2 Distributional Effects of Macroeconomic Shocks

Having considered the *aggregate* effects of macroeconomic policies, I now turn my attention to the *distributional* effects. To do this, I consider a policymaker who wants to stabilize the economy in the face of economic shocks. One tool the policymaker could use is monetary policy. Other tools are fiscal policy, i.e., changes in government consumption or government transfers. Any of these tools can stabilize aggregate demand. However, they might have different distributional effects, i.e., affect households differently. To study this, I ask the following question: For each \$100 generated by a policy, how much goes to the top x%? This question helps clarify whether policies favor the poor or the rich.

In Table 3, I ask exactly this question for the top 0.1%. I ask the question in both models for three different shocks: Monetary policy, transfers, and government consumption. The table shows a striking result: In the model with heterogeneous returns, around 11% of all income generated by monetary policy goes to the top 0.1%.[22] These numbers are much lower for transfers and government consumption, at around 1% and 3%, respectively. Furthermore, the number is much larger than in the standard HANK model, where all policies modestly favor the top 0.1% with less than 2% of income going to the top 0.1%.

	Standard HANK	Heterogeneous r^a
Transfers	0.1%	1.1%
Government consumption	0.6%	2.9%
Monetary policy	1.8%	10.7%

Table 3: Shares of Income Going to the Top 0.1% In Response to Shocks

Note: The table shows the shares of income going to the top 0.1% on impact in response to three different policies in two models.

Next, I generalize Table 3 from the top 0.1% to any point in the wealth distribution in Figure 15. The figure shows that transfers are by far the most equal policy: It benefits households across the income distribution fairly equally, with the bottom x% getting almost the corresponding share of the increase in income. As an example,

22. Appendix E.5 shows the difference between the models for the whole distribution.

the bottom 80% get around 70% of the income generated by the policy. This is not surprising given that transfers are lump-sum in the model, benefiting households equally. Transfers are not completely equal only due to general equilibrium effects.[23] The second most equal policy is government consumption. Here, the bottom 80% get around 40% of the income generated. Consistent with Table 3, monetary policy is by far the least equal policy, with the bottom 80% getting essentially only around 15% of income. Note that around a third of all households are borrowing-constrained and therefore have zero assets. These 1/3 of households are the flat segment in Figure 15, which also explains why the policies lie above the 45-degree line at the bottom.[24]

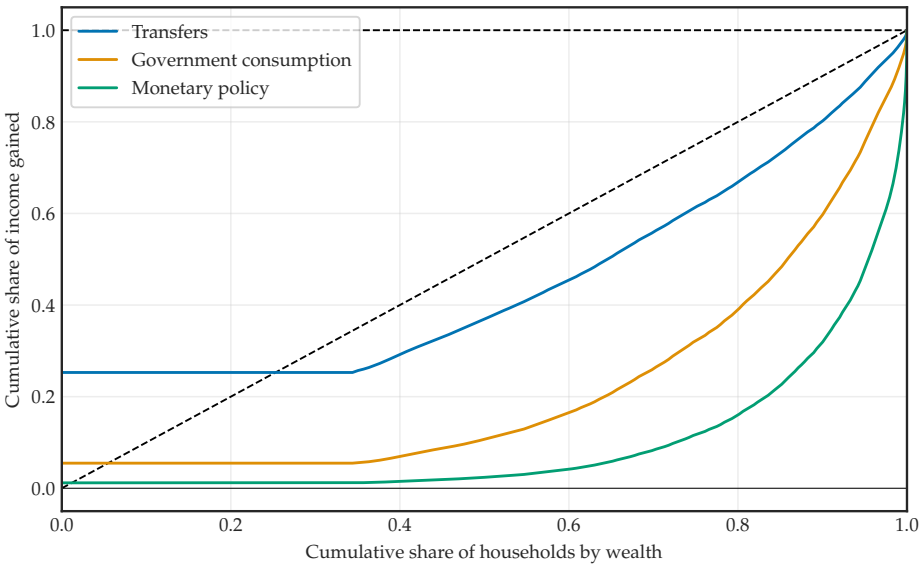


Figure 15: Shares of Income Going to Different Households

Note: The figure shows the composition of income generated by different policies on impact along the wealth distribution. The plot is smoothed using a Savitzky-Golay filter to smooth out kinks in the solution due to the discretization.

Why is this? The key thing to understand is that monetary policy implies a positive revaluation of wealth, $dr_0^a > 0$, as argued previously. How does this translate into income for different households? Note that at the top of the wealth distribution, income is essentially only capital income, $a_{i,-1} r_{i0}^a$. This approximation certainly holds well for the top 0.1%. The gains from higher asset prices are much higher at the top

23. For instance, higher transfers create a boom, which causes dividends to increase since they are pro-cyclical. Higher dividends mostly accrue to the rich firm owners.

24. Figure 15 shows the share of income going to households with $a_{i,ss} \leq \mathrm{P}x$, where $\mathrm{P}x$ is the x'th percentile of wealth. Since $\mathrm{P}x = 0$ for all $x < 1/3$, the bottom $x < 1/3$ of households in the figure is always the bottom $1/3$. Alternatively, the figure could use a strict inequality, $a_{i,ss} < \mathrm{P}x$. Then, the bottom $x < 1/3$ would be the bottom 0 with a discrete jump around $1/3$.

of the wealth distribution for two reasons in the model with heterogeneous returns: Because top wealth is much larger, and because the pass-through of returns is larger. This is why monetary policy benefits the ultra-rich much more in the model with heterogeneous returns than the standard HANK model, i.e., the third row in Table 3.

What explains the first two rows in Table 3, i.e., why do transfers and government consumption benefit the ultra-rich much less than monetary policy? This is quite simple: Even if the policies induce the same change in output, monetary policy inflates asset prices more due to lower discount rates, benefiting the ultra-rich more.

One way to see the importance of capital income is in Table 4, which splits the aggregate change of income into the three sources: Labor income, z_{it}, capital income, $r^a_{it} a_{it-1}$, and transfer income, T_t. Table 4 shows each of the three income sources for each of the three policies. The table shows that transfers mostly generate transfer income, unsurprisingly. In contrast, government consumption mostly generates labor income: This is because the government buys goods, and firms pay households to produce these goods. Furthermore, monetary policy mainly generates capital income. Across the wealth distribution, the income sources are more or less the same in response to both fiscal policy shocks, while the income sources are uneven in response to monetary policy: The richest households get capital income mostly, while the poorer households get labor income, cf. Appendix E.6.

	Capital share	Labor share	Transfer share
Transfers	8%	26%	67%
Government consumption	22%	78%	0%
Monetary policy	83%	17%	0%

Table 4: Income Composition of Different Policies

Note: The table shows the composition of income generated by different policies on impact.

Are the distributional effects I find consistent with the empirical literature? I find that this is the case. For instance, Andersen et al. (2023) find that the gains of easy monetary policy are increasing in income. I confirm that this is also the case in my model, cf. Appendix E.7.

6.3 Robustness

Let me now consider the robustness of the results. To do so, I present a series of model changes and their implications for the model. In particular, I consider the following changes to the model: Adding sticky prices, considering long-term debt, considering nominal debt, adding capital to the model, and parametrizing the model differently. All the changes are spelled out in Appendix E.8. I only review the results here.

The results of all model changes are shown in Table 5. As the table shows, the results are very robust to all changes. In particular, the increase in output remains in a narrow band. The same is the case for the top 0.1% income share of monetary policy, with one slight exception: When more liquidity is in the form of bonds, the top 0.1% income share drops slightly. This is because there is no revaluation of bonds, as government bonds guarantee some real return by assumption, so capital income from holding them does not increase as with firm equity. Despite the implausibly large value of government bonds, however, the top 0.1% income share remains large.

Model	Output increase	Top 0.1% income share
Baseline	1.45%	10.7%
Capital	1.30%	14.1%
Long-term debt	1.50%	10.9%
Nominal debt	1.46%	10.6%
Sticky prices	1.44%	10.6%
More flexible wages	1.45%	10.7%
More bonds	1.47%	9.1%
Net borrowing	1.37%	12.2%

Table 5: Robustness

Note: The table shows the robustness of the results to various changes to the model.

7 Welfare

7.1 The Equivalent Variation of Shocks

So far, I have considered how households are affected differently in terms of *income*. While understanding the income effects is intuitive and a step to understanding the welfare effects, the income effects do not necessarily translate directly to *welfare*. For this reason, I now focus on the welfare effects of shocks across the distribution.

One of the main issues with considering income is how capital income is affected. In particular, capital income rises on impact but drops afterward. The importance of considering welfare instead of asset price revaluations is emphasized by Fagereng et al. (2025). Whether households would prefer not having their income change is a complex question. To handle this, I consider a welfare measure of the shock. In particular, I largely follow the approach from Bardóczy and Velásquez-Giraldo (2024). Let $V_{ss}(e_i, a_i)$ denote the value function at state (e_i, a_i) in steady state. Consider then an unexpected aggregate shock. Denote the value function after this shock by $V^*(e_i, a_i)$. The equivalent variation, ev_i, then solves[25]

$$V^*(e_i, a_i) = V_{ss}\left(e_i, a_i + \frac{\text{ev}_i}{1 + (1-\tau)r_i^a}\right).$$

This answers the question: What is the transfer that would make the household indifferent between facing the shock and not facing the shock? Positive values, $\text{ev}_i > 0$, indicate that the household needs to be compensated not to face the shock, i.e., that the household likes the shock, and vice versa for negative values.

Figure 16 re-creates Figure 15 showing income shares by now showing equivalent variations instead. The figure focuses on monetary policy and transfers, as government consumption is wasteful and therefore is hard to interpret when looking at equivalent variations. The figure has essentially the same takeaways as Figure 15 based on income: The gains of monetary policy are significantly less evenly distributed than the gains of fiscal policy. However, the equivalent variation gains of all shocks are more evenly distributed compared to the income gains. But the ranking and relative difference between shocks remain the same. If anything, the *differences* in how unevenly shocks are distributed are even more pronounced when looking at equivalent variations than with income.

25. I divide by $1 + (1-\tau)r_i^a$ since increasing wealth by x increases cash-on-hand by $(1 + (1-\tau)r_i^a)x$.

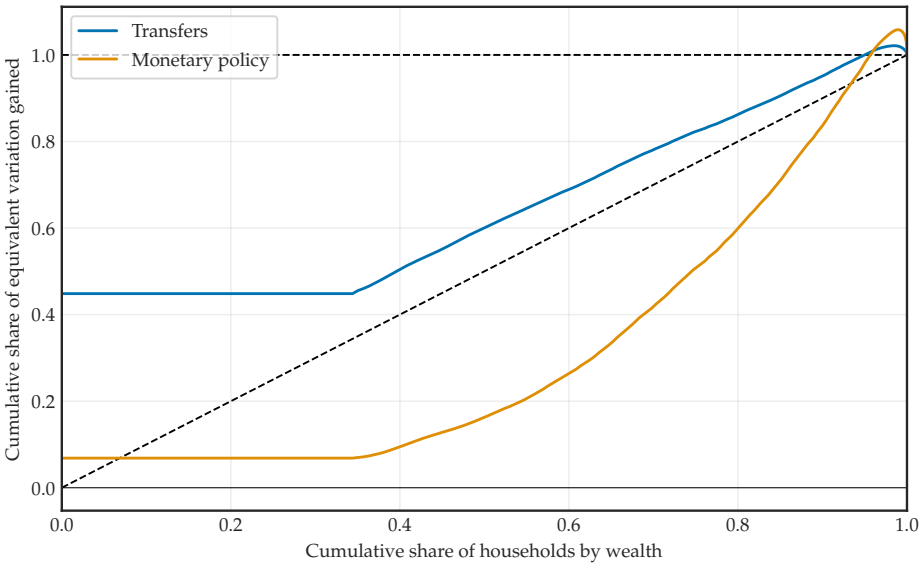


Figure 16: Shares of Equivalent Variations Going to Different Households

Note: The figure shows the composition of equivalent variations generated by different policies on impact along the wealth distribution. The plot is smoothed using a Savitzky-Golay filter to smooth out kinks in the solution due to the discretization.

One disadvantage of looking at equivalent variations instead of income is that the underlying drivers are less transparent. To get at this, I decompose the equivalent variation gains observed in Figure 16 in Appendix F.1. This shows that the gains of monetary policy at the top largely come from the increase in asset prices.

7.2 Policy Implications

I round out the paper by considering the effects of different policies on welfare. To do so, I do as follows. Over the business cycle, the economy is hit by shocks. A policymaker can stabilize business cycles by following policy rules. The classic example is a monetary policy rule for the interest rate, but policymakers could also use a rule for fiscal policy. Motivated by the different effects of shocks in Section 6, I study rules for both monetary policy and fiscal policy. I then ask the question: Do different households prefer if stabilization is conducted using different rules?

To do so, I first add a business cycle to the economy. In particular, I introduce three shocks to the model: A markup shock, a risk premium shock, and a Phillips curve shock. I assume that all three shocks follow AR(1) processes with IID normal innovations. I then estimate these shocks to fit actual US business cycles in the data. All the details are given in Appendix F.2, but the key point is just to have some variation in business cycles, which approximates empirical business cycles, for the

policy to stabilize.

With the shocks estimated, I simulate the economy getting hit by the three shocks for 1,000,000 years. The goal of the exercise is then to compare different policy rules used to stabilize the business cycles. Given the results showing that the gains from monetary policy are the least equally distributed, while transfers are the most equally distributed, one might think that rich households prefer using monetary policy to stabilize business cycles, while poor households prefer fiscal policy. But this is not the case because preferences are symmetric: Just as much as rich households like stabilizing recessions by lowering the interest rate, they dislike stabilizing booms by increasing the interest rate. For this reason, I instead propose to look at asymmetric policies.

In particular, I consider two different asymmetric policies. The first is where fiscal policy in the form of transfers is used to stabilize shocks that create recessions and monetary policy to stabilize shocks that create booms, while the second is the exact opposite. Appendix F.3 gives the details on the simulation. Crucially, both asymmetric policies are set such that the business cycle is perfectly stabilized at all points in time, i.e., $dY_t = dC_t = 0$. This allows me to compare the distributional effects holding fixed aggregate stabilization. I then compute welfare—as presented above—for both poor and rich households in these two scenarios to compare which households prefer which scenarios. Here, the "poor" are defined as the bottom 90% by wealth (holding 23% of all wealth), while the "rich" are the top 10% (holding the remaining 77%). The results are given in Table 6. In particular, Table 6 reports average equivalent variations for both groups in percent of GDP compared to the steady state with no business cycles.

Policy	Poor	Rich
FP in recessions + MP in booms	1.22%	−0.24%
MP in recessions + FP in booms	−1.22%	0.24%

Table 6: Welfare Gains of Asymmetric Policies

Note: The table shows equivalent variation welfare gains in percent of GDP from two different asymmetric policies compared to steady state. Fiscal policy refers to lump-sum transfers, while monetary policy refers to the real interest rate.

Table 6 shows that poor households clearly prefer the asymmetric policy using fiscal policy in recessions and monetary policy in booms. On the other hand, rich

households prefer the opposite. By construction, the rows are the same with opposite sign. This is not surprising given the distributional effects of the different policies. Appendix F.4 shows that the results in Table 6 are robust.

Interestingly, one can argue that the asymmetric policy using transfers in recessions but not in booms is similar to what policymakers have actually done in the last decades in the US. In particular, the federal surplus has displayed an asymmetry, dropping sharply in recessions without increasing similarly in booms. Examples of such policies include various rounds of transfers, including tax rebate checks during the 2001 recession, similar checks in 2008 during the financial crisis, and finally three rounds of "stimulus checks" in 2020 and 2021 during the Covid recession. At the same time as the asymmetric policy has been conducted, monetary policy has been more symmetrical, with rates decreasing in recessions and increasing in booms. Appendix F.5 shows examples of the simulated time series in the model.

Let me end this section with two important notes. First, this analysis says nothing about which policy is optimal. Indeed, stabilizing output and consumption in response to all shocks is generally not optimal. Saying something about optimal policy would require taking a normative stance on how to weight the welfare of different households and, thus, to what degree the policymaker cares about inequality. For instance, some papers use a utilitarian objective where all households are weighted equally (Dávila and Schaab 2023, Le Grand and Ragot 2023), while other papers weight rich or poor households relatively more or less (Bhandari et al. 2021, McKay and Wolf 2023b). Instead, my analysis simply positively identifies which policies different households would prefer, side-stepping the normative issues. Thus, my results can serve as the basis for informing a normative analysis.

Second, my analysis does not take a stance on how monetary policy or fiscal policy should be conducted. Instead, it simply compares two rules that achieve the same stabilization of aggregate output and consumption. In this sense, my paper is different from Gornemann et al. (2016), who compare how monetary policy should be conducted, i.e., how hawkish or dovish it should be.

8 Conclusion

I study the distributional effects of monetary policy when returns are heterogeneous in an otherwise standard HANK model. I do so by constructing a dataset of heterogeneous returns across US households. I have three main takeaways. First, I find that the model matches micro distributions of returns, wealth, and income. Crucially, the model matches the concentration of wealth at the top and the pass-through of aggregate to households' returns.

Second, I study the distributional effects of monetary policy. I find that income gains from expansionary monetary policy disproportionately benefit the ultra-rich: The top 0.1% take 11% of the income increase, more than 100 times their population share and an order of magnitude more than in standard HANK models. This is because monetary policy mainly increases capital income, which mostly goes to the rich. On the other hand, fiscal policy is much more equally distributed.

Third, I study the policy implications. I do this by considering asymmetric policy over the business cycle. For instance, policymakers can ease fiscal policy in recessions and tighten monetary policy in booms. I find that poor households prefer this. However, there is disagreement about the policy: Rich households prefer the exact opposite, using easy monetary policy in recessions and tight fiscal policy in booms.

My paper thus constitutes new evidence on the distributional effects of monetary policy. Policymakers who care about distributional effects should consider this when designing stabilization policies.

References

Alves, Felipe, Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. 2020. "A Further Look at the Propagation of Monetary Policy Shocks in HANK." *Journal of Money, Credit and Banking* 52, no. S2 (December): 521–559.

Andersen, Asger Lau, Niels Johannesen, Mia Jørgensen, and José-Luis Peydró. 2023. "Monetary Policy and Inequality." *The Journal of Finance* 78, no. 5 (October): 2945–2989.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109, no. 6 (June): 2333–2367.

Auclert, Adrien, Bence Bardóczy, and Matthew Rognlie. 2023. "MPCs, MPEs, and Multipliers: A Trilemma for New Keynesian Models." *Review of Economics and Statistics* 105, no. 3 (May): 700–712.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. 2021. "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models." *Econometrica* 89 (5): 2375–2408.

Auclert, Adrien, and Matthew Rognlie. 2018. "Inequality and Aggregate Demand." *NBER Working Paper 24280.*

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2020. "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model." *NBER Working Paper 26647.*

———. 2024a. "Fiscal and Monetary Policy with Heterogeneous Agents." *NBER Working Paper 32991.*

———. 2024b. "The Intertemporal Keynesian Cross." *Journal of Political Economy* 132 (12): 4068–4121.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2020. "Rich Pickings? Risk, Return, and Skill in Household Wealth." *American Economic Review* 110, no. 9 (September): 2703–2747.

Bardóczy, Bence, and Mateo Velásquez-Giraldo. 2024. "HANK Comes of Age." *Finance and Economics Discussion Series,* nos. 2024-052 (July): 1–64.

Barro, Robert J., and Charles J. Redlick. 2011. "Macroeconomic Effects From Government Purchases and Taxes." *The Quarterly Journal of Economics* 126, no. 1 (February): 51–102.

Bayer, Christian, Benjamin Born, and Ralph Luetticke. 2024. "Shocks, Frictions, and Inequality in US Business Cycles." *American Economic Review* 114, no. 5 (May): 1211–1247.

Benhabib, Jess, Alberto Bisin, and Mi Luo. 2017. "Earnings Inequality and Other Determinants of Wealth Inequality." *American Economic Review* 107, no. 5 (May): 593–597.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu. 2011. "The Distribution of Wealth and Fiscal Policy in Economies With Finitely Lived Agents." *Econometrica* 79 (1): 123–157.

———. 2015. "The wealth distribution in Bewley economies with capital income risk." *Journal of Economic Theory* 159 (September): 489–515.

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent. 2021. "Inequality, Business Cycles, and Monetary-Fiscal Policy." *Econometrica* 89 (6): 2559–2599.

Bilbiie, Florin, Sigurd Galaasen, Refet Gürkaynak, Mathis Maehlum, and Krisztina Molnar. 2025. "HANKSSON." *CEPR Discussion Paper No. 20090.*

Boppart, Timo, Per Krusell, and Kurt Mitman. 2018. "Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative." *Journal of Economic Dynamics and Control* 89 (April): 68–92.

Broer, Tobias, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Öberg. 2020. "The New Keynesian Transmission Mechanism: A Heterogeneous-Agent Perspective." *The Review of Economic Studies* 87, no. 1 (January): 77–101.

Carroll, Christopher D. 2006. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics Letters* 91, no. 3 (June): 312–320.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary policy and inequality." *Journal of Monetary Economics* 88 (June): 70–89.

Daminato, Claudio, and Luigi Pistaferri. 2024. "Returns Heterogeneity and Consumption Inequality Over the Life Cycle." *NBER Working Paper 32490.*

Dávila, Eduardo, and Andreas Schaab. 2023. "Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy." *NBER Working Paper 30961.*

de Ferra, Sergio, Kurt Mitman, and Federica Romei. 2020. "Household heterogeneity and the transmission of foreign shocks." *Journal of International Economics* 124 (May): 103303.

Druedahl, Jeppe, Søren Hove Ravn, Laura Sunder-Plassmann, Jacob Marott Sundram, and Nicolai Waldstrøm. 2025. "Fiscal Multipliers in Small Open Economies With Heterogeneous Households." *IMF Economic Review.*

Fagereng, Andreas, Matthieu Gomez, Émilien Gouin-Bonenfant, Martin Holm, Benjamin Moll, and Gisle Natvik. 2025. "Asset-Price Redistribution." *Journal of Political Economy* 133 (11): 3494–3549.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and Persistence in Returns to Wealth." *Econometrica* 88 (1): 115–170.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. "MPC Heterogeneity and Household Balance Sheets." *American Economic Journal: Macroeconomics* 13, no. 4 (October): 1–54.

Floden, Martin, and Jesper Lindé. 2001. "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" *Review of Economic Dynamics* 4, no. 2 (April): 406–437.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. "The Dynamics of Inequality." *Econometrica* 84 (6): 2071–2111.

Gaillard, Alexandre, Philipp Wangner, Christian Hellwig, and Nicolas Werquin. 2023. "Consumption, Wealth, and Income Inequality: A Tale of Tails." *SSRN Working Paper 4636704.*

Gornemann, Nils, Keith Kuester, and Makoto Nakajima. 2016. "Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy." *International Finance Discussion Paper* 2016, no. 1167 (May): 1–40.

Gouin-Bonenfant, Émilien, and Alexis Akira Toda. 2023. "Pareto extrapolation: An analytical framework for studying tail inequality." *Quantitative Economics* 14 (1): 201–233.

Guvenen, Fatih, Gueorgui Kambourov, Burhan Kuruscu, Sergio Ocampo, and Daphne Chen. 2023. "Use It or Lose It: Efficiency and Redistributional Effects of Wealth Taxation." *The Quarterly Journal of Economics* 138, no. 2 (April): 835–894.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" *Econometrica* 89 (5): 2303–2339.

Guvenen, Fatih, Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo. 2017. "Worker Betas: Five Facts about Systematic Earnings Risk." *American Economic Review* 107, no. 5 (May): 398–403.

Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman. 2019. "The Fiscal Multiplier." *NBER Working Paper 25571.*

Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek. 2021. "The Transmission of Monetary Policy under the Microscope." *Journal of Political Economy* 129, no. 10 (October): 2861–2904.

Jappelli, Tullio, and Luigi Pistaferri. 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics* 6, no. 4 (October): 107–136.

Jones, Charles I., and Jihee Kim. 2018. "A Schumpeterian Model of Top Income Inequality." *Journal of Political Economy* 126, no. 5 (October): 1785–1826.

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor. 2019. "The Rate of Return on Everything, 1870–2015." *The Quarterly Journal of Economics* 134, no. 3 (August): 1225–1298.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108, no. 3 (March): 697–743.

Kaplan, Greg, and Giovanni Violante. 2018. "Microeconomic Heterogeneity and Macroeconomic Shocks." *Journal of Economic Perspectives* 32, no. 3 (August): 167–194.

———. 2022. "The Marginal Propensity to Consume in Heterogeneous Agent Models." *NBER Working Paper 30013.*

Kopecky, Karen A., and Richard M.H. Suen. 2010. "Finite state Markov-chain approximations to highly persistent processes." *Review of Economic Dynamics* 13, no. 3 (July): 701–714.

Kuhn, Moritz, and Jose-Victor Rios-Rull. 2016. "2013 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomic Modelers."

Le Grand, François, and Xavier Ragot. 2023. "Optimal policies with heterogeneous agents: Truncation and transitions." *Journal of Economic Dynamics and Control* 156 (November): 104737.

McKay, Alisdair, and Christian K. Wolf. 2023a. "Monetary Policy and Inequality." *Journal of Economic Perspectives* 37, no. 1 (February): 121–144.

———. 2023b. "Optimal Policy Rules in HANK." *Working Paper.*

Menzio, Guido, and Saverio Spinella. 2025. "A Quantitative Theory of Heterogeneous Returns to Wealth." *NBER Working Paper 33868.*

Nielsen, Roger. 2006. *An Introduction to Copulas.* Springer Series in Statistics. New York, NY: Springer New York.

Orchard, Jacob D, Valerie A Ramey, and Johannes F Wieland. 2025. "Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates." *The Quarterly Journal of Economics* 140, no. 3 (July): 2001–2052.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. "Distributional National Accounts: Methods and Estimates for the United States." *The Quarterly Journal of Economics* 133, no. 2 (May): 553–609.

Ravn, Morten O, and Vincent Sterk. 2021. "Macroeconomic Fluctuations with HANK & SAM: an Analytical Approach." *Journal of the European Economic Association* 19, no. 2 (April): 1162–1202.

Rouwenhorst, K. Geert. 1995. "Asset Pricing Implications of Equilibrium Business Cycle Models."

Smith, Matthew, Owen Zidar, and Eric Zwick. 2022. "Top Wealth in America: New Estimates Under Heterogeneous Returns." *The Quarterly Journal of Economics* 138, no. 1 (December): 515–573.

Snudden, Stephen. 2021. "Leverage and Rate of Return Heterogeneity among U.S. Households." *Working Paper.*

Vermeulen, Philip. 2018. "How Fat is the Top Tail of the Wealth Distribution?" *Review of Income and Wealth* 64, no. 2 (June): 357–387.

Werning, Iván. 2015. "Incomplete Markets and Aggregate Demand." *NBER Working Paper 21448.*

Xavier, Inês. 2021. "Wealth Inequality in the US: the Role of Heterogeneous Returns." *SSRN Working Paper 3915439.*

Appendix

A Appendix to Section 2

A.1 Wealth Concentration

Table A.1 shows top wealth shares in the data and selected HANK models in the literature. I have selected these HANK models because of their prominence in the literature, their emphasis on matching the distribution of wealth, and their reporting of top wealth shares. As such, the models in Table A.1 provide the best case of matching the wealth concentration in standard HANK models, with most HANK models featuring *less* wealth concentration.

Table A.1 shows that all four models understate the concentration of wealth at the top. In particular, Bayer et al. (2024) (BBL) only report the top 10% share, while Auclert et al. (2020) (ARS) only report top 10% and 5% shares, understating these compared to the data. Their fit to the top 1% or 0.1% shares is then almost surely worse. McKay and Wolf (2023b) (MW) fit the top 5% share, but understate the top 1% share and do not report the top 0.1% share. Kaplan et al. (2018) (KMV) clearly do the best: They fit the top 10%, 5% share, and top 1%. Their fit only fails at the very top: They do not fit the top 0.1% share, something they are open about in the paper. However, it is worth noting that they are also the only ones to report the top 0.1% share.

The models in Table A.1 do not report their Pareto tail index. This is natural, as the models do not feature a Pareto tail and thus the Pareto tail index is not defined in the models.

	Data	KMV	ARS	MW	BBL
Top 10% share	76%	82%	70%	82%	67%
Top 5% share	65%	69%	58%	66%	—
Top 1% share	37%	38%	—	27%	—
Top 0.1% share	14%	7%	—	—	—
Pareto tail index	1.52	—	—	—	—

Table A.1: Wealth Concentration in the Data and HANK Models

Note: The table shows wealth shares and the Pareto tail index of wealth in the US data and selected models. The data is the 2019 SCF, except the Pareto tail index, which is from Vermeulen (2018). "KMV" is Kaplan et al. (2018), "ARS" is Auclert et al. (2020), "MW" is McKay and Wolf (2023b), and BBL is Bayer et al. (2024). The numbers for ARS are for the illiquid wealth distribution.

A.2 Survey of Consumer Finances

I use data from the 2019 SCF. The data is from Kuhn and Rios-Rull (2016), which is updated to 2019 online. They provide data on earnings, income, and wealth directly. They also provide data on transfers as a share of income, which I convert to USD using the income data. I do the same for "other" income.

In the model, income is separated into earnings, capital income, and transfers. There are two issues to overcome in order to make the data consistent with the model: (1) What to consider as capital income, and (2) how to attribute "other" income. For (1), I compute capital income as the part of income not due to earnings, transfers, or other income. Regarding (2), I distribute other income to the three remaining components (earnings, transfers, and capital income). The result is a variable for income identical to the one in the data provided by Kuhn and Rios-Rull (2016), but which is made up of three components: (1) Earnings, (2) capital income, and (3) transfers.

A.3 Piketty-Saez-Zucman

I use the micro-files for the distributional national accounts for the US from Piketty et al. (2018). The data is constructed based on tax, survey, and national account data. I use the variable *hweal* ("net personal wealth") as my measure of wealth and the variable *pkinc* ("personal pre-tax capital income") as my measure of capital income.

As with the PSID, I drop households older than 65.

A.4 Panel Study of Income Dynamics

In this section, I discuss how I use data from the PSID to construct a dataset of heterogeneous returns across US households. The construction largely follows Snudden (2021).

A.4.1 Data Structure

In addition to the SCF, I use panel data from the PSID. The panel is biennial, starting in 1999 and ending in 2019. Crucially, the unit of time for flow variables is still years. Thus, the panel can be thought of as annual with missing data.

The variables I use can be considered as being part of two categories: Wealth *stocks* and income *flows*. For a survey conducted in year t, respondents are asked about their *stock* of wealth at that point in time, i.e., year t, and about income flows during the previous year, i.e., year $t-1$, see Figure A.1. For this reason, both the numerator and denominator in eq. (2) are observed, though not in the same survey. Thus, I can construct returns for year $t-1$ for each survey year t, but not t itself. Since surveys were conducted in 1999, 2001, ..., 2019, I can construct returns for 2000, 2002, ..., 2018.[26]

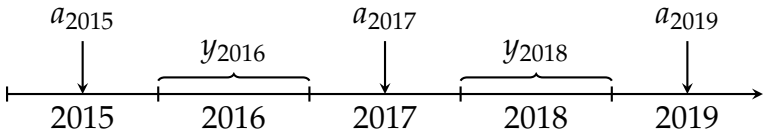


Figure A.1: Example of Survey Timeline

A.4.2 Sample Restrictions

I restrict the age of the head to be 18–70. I drop the SEO sample consisting of oversampled low-income households. I restrict attention to households that have the same head. I trim the sample by dropping returns below the 0.5th and above the 99.5th percentiles. When computing returns using the formula in eq. (2), I require that households have $a_{it-1} > 1000$ USD to ensure that I am not dividing by a number close to zero. Returns are adjusted for consumer price inflation.

26. I lose the first year due to the lagged wealth in eq. (2).

A.4.3 Capital Income

I now discuss how I compute the capital income from the eight sources. First, I discuss income from trust funds and royalties, interest income, and dividend income. Respondents are asked directly about these three in the survey, so I simply use the responses. I discuss the remaining five in turn.

- *Primary Housing.* Income from primary housing can be split into two parts. The first part is rental income, which is reported directly in the survey, but is not attributed to primary or other housing. I attribute the rental income to primary housing if the household is a homeowner and does not own other real estate. Otherwise, I attribute it to other housing.

  The second part is capital gains. Capital gains are computed as

  $$\text{capital gains} = \frac{\text{price change} - \text{improvements}}{2}.$$

  The price change is the change in the price of the house, which depends on whether the house was sold or not. If the house was sold, the selling price is used. If the house was not sold, the self-assessed value is used. Both are net of an 8% commission. Improvements are reported directly in the survey but are not attributed to primary or other housing. I attribute improvements to primary housing if the household is a homeowner, and otherwise, I attribute them to other housing.

- *Other Housing.* Income from other housing can be split into the same two parts as primary housing. The first part is rental income, which is reported directly in the survey, but is not attributed to primary or other housing. I attribute it to what was presented when discussing primary housing.

  The second part is capital gains. Capital gains are computed as

  $$\text{capital gains} = \frac{\text{price change} - \text{improvements} - \text{net investment}}{2}.$$

  The price change is simply the change in the price, while net investments are the difference between the price of real estate sold and the price of real estate bought. Improvements are as discussed with primary housing.

- *Businesses.* Business income is split into two types: Realized and unrealized capital gains. The realized part is the sum of the income associated with owning

the business—reported directly in the survey—and the income associated with owning the farm.[27]

- *Stocks.* Income from stocks is given by

$$x_{it}^{\text{stocks}} = \frac{\Delta a_{it}^{\text{stocks}} - f_{it}}{2},$$

where a_{it}^{stocks} is the value of the stocks and f_{it} is the net investment into stocks.

- *Other.* In the survey, the respondent is asked for the total capital income of other members of the family unit. I include this as *other* capital income.[28]

A.4.4 Wealth

The value of wealth can be split into two types in the survey: Assets for which net investment *is* reported, and assets for which net investment is *not* reported. This distinction is important because it matters for how returns are computed. In particular, it matters for the measurement of the denominator in eq. (2). When net investments are reported, I simply use the lagged value of wealth in the denominator. When net investments are not reported, I use an average of the lagged value of wealth and the contemporaneous value of wealth. Let me start by discussing the assets where net investment is reported.

- *Primary housing.* Respondents are asked about the value of their house. I report the value net of 8% commission.

- *Other housing.* Respondents are asked about the value of their real estate, i.e., how much it would sell for. Values are kept consistent across waves and reconciled with reported transactions around survey question changes, so the measure reflects a market-value concept across the panel.

- *Stocks.* Respondents are asked about the value of their stocks if they paid off everything owed on them.

27. Farm income is not split into capital and labor, so I do this. If farm income is negative, I attribute all of it to capital income. If it is positive, I attribute half to capital income.

28. Before 2005, respondents were only asked about the *total* income of other members of the family unit, not how much of it is labor income. I attribute this income as labor income before 2005, except if it is negative, in which case I attribute it as capital income.

- *Businesses.* Respondents are asked about the value of their farm/business, i.e., how much it would sell for. Reported values are reconciled with buys/sells and changes in reported net worth between waves (including liability repayments when reported), so changes reflect both transactions and valuation movements.

Let me then discuss the assets where net investment is *not* reported.

- *Private annuities or IRAs.* Respondents are asked about the value of their annuities/IRAs.
- *Checking/savings accounts.* Respondents are asked about the value of their checking/savings accounts.
- *Vehicles.* Respondents are asked about the value of their vehicles if they paid off everything owed on them.
- *Other assets.* Respondents are asked about the value of other assets if they paid off everything owed on them.

A.4.5 Descriptive Statistics

Table A.2 shows key descriptive statistics for the PSID data.

	Mean	Std. dev.	P5	P50	P95
Return	1.2%	13.3%	−12.1%	−1.7%	21.8%
Assets (USD)	318	954	2	138	1149
— Primary house	141	200	0	92	460
— Other house	32	601	0	0	150
— Business	32	393	0	0	10
— Stocks	29	209	0	0	100
— IRAs	41	175	0	0	220
— Savings	20	85	0	3	83
— Vehicles	16	21	0	10	50
— Other	9	85	0	0	25
Age (years)	43	13	25	42	65

Table A.2: Descriptive Statistics for Heterogeneous Returns and Wealth in the PSID

Note: The values of wealth are in thousands of nominal USD. The age is in years.

A.4.6 Explaining Returns With Portfolio Shares

Consider splitting household wealth into N assets indexed by $j = 1, \ldots, N$ such that total capital income and wealth are the sum over these assets:

$$x_{it} = \sum_{j=1}^{N} x_{it}^{j} \quad \text{and} \quad a_{it} = \sum_{j=1}^{N} a_{it}^{j}.$$

Consider the case where households only earn different returns because they hold different assets, but each asset gives the same return, i.e., $x_{it}^{j} = r_t^j a_{it-1}^{j}$. In this case, the total return for a household is just a weighted average of households' returns,

$$r_{it} = \frac{x_{it}}{a_{it-1}} = \sum_{j=1}^{N} \frac{a_{it-1}^{j}}{a_{it-1}} \frac{x_{it}^{j}}{a_{it-1}^{j}} = \sum_{j=1}^{N} \omega_{it-1}^{j} r_t^j, \tag{15}$$

where the weights are $\omega_{it}^{j} = a_{it}^{j}/a_{it}$. An immediate implication is that a regression of households' returns on households' weights should yield an R^2—coefficient of determination—of exactly 1 *within each year*.

Figure A.2 reports the R^2 of such regressions.[29] The R^2 is in the range of 0.01–0.06, far away from 1. Even significant measurement error cannot explain this, suggesting that household returns are idiosyncratic also within asset categories and years. This is consistent with Fagereng et al. (2020), who find that returns are heterogeneous also within narrow asset classes.

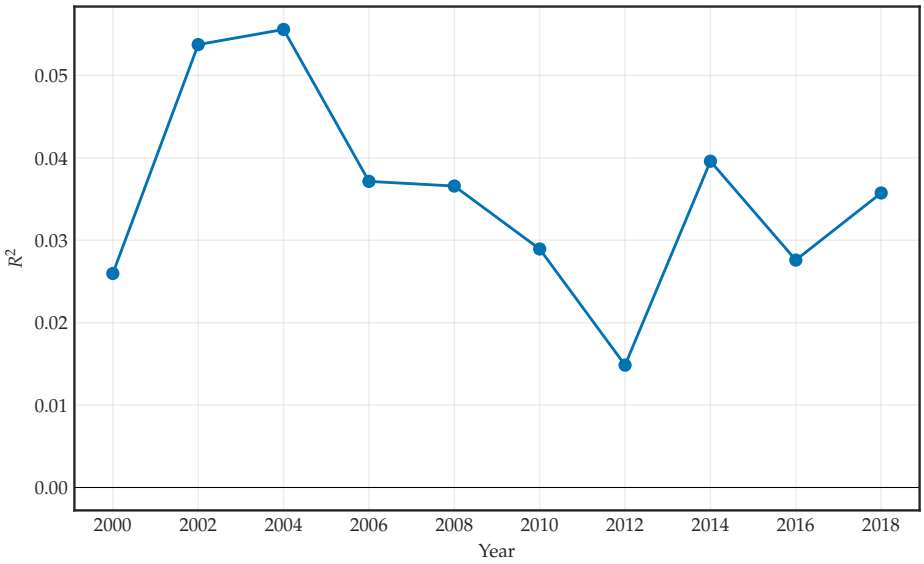


Figure A.2: R^2 of Regression of Returns on Portfolio Shares by Year

Note: The figure shows the coefficient of determination, i.e., R^2, of a regression of household-level returns on their portfolio shares by year.

29. I include a constant in the regression even though eq. (15) suggests that this should be 0. This should only bias upwards the R^2.

A.4.7 Return Pass-Through Robustness

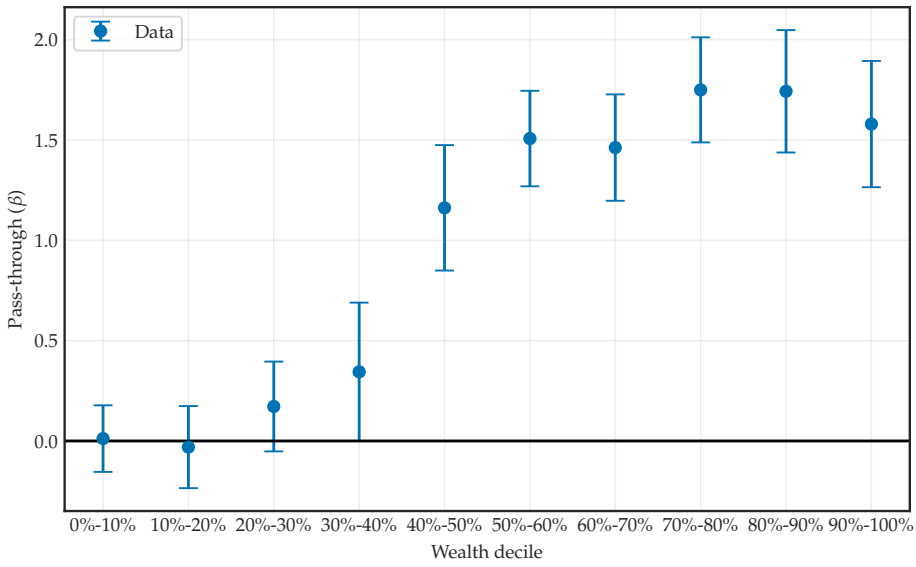


Figure A.3: Pass-Through of Average to Households' Returns by Wealth With Standard Errors Clustered by Household

Note: See Figure 4. The only difference is that the standard errors are clustered by household. Confidence intervals are at the 95% level.

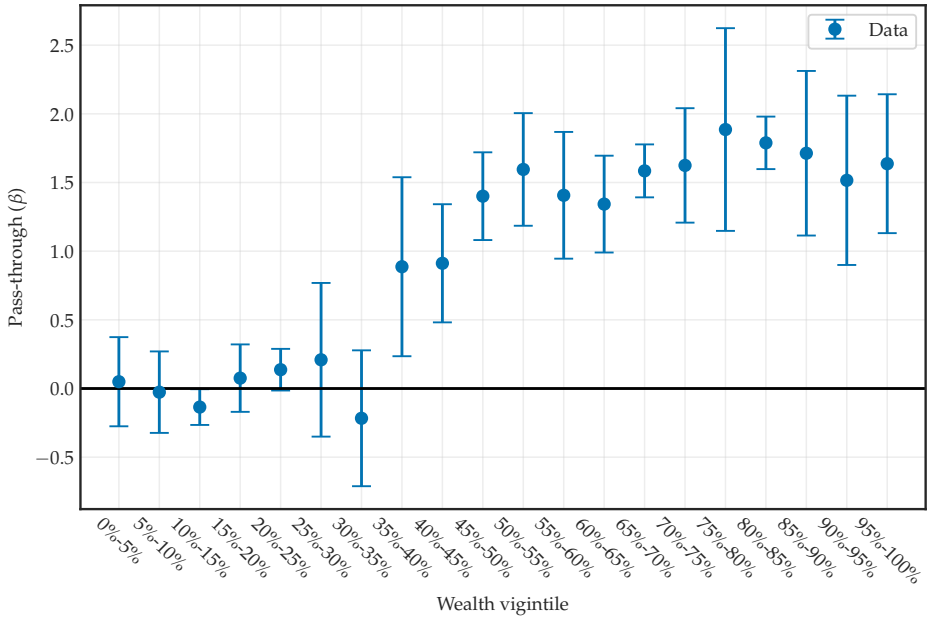


Figure A.4: Pass-Through of Average to Households' Returns by Wealth Vigintiles

Note: Figure 4 for vigintiles (groups of 20) instead of deciles (groups of 10). Confidence intervals are at the 95% level.

A.4.8 Robustness to Asset Types

In this appendix, I consider removing some asset types from the data. In particular, I consider three alternative return variables: (i) Without housing, (ii) without vehicles, and (iii) financial assets only. In Table A.3, I report the average return and standard deviation of returns for these variables. These alternative returns generally have a larger standard deviation, particularly when looking at financial variables only.

	Mean	Std. dev.
Everything	1.2%	13.3%
W/o housing	1.1%	15.1%
W/o vehicles	2.6%	17.4%
Financial	4.1%	26.9%

Table A.3: Alternative Returns Variables

Note: The table shows the average return and the standard deviation of returns with alternative returns variables.

In Figure A.5, I consider two key figures in the alternative return variables. The left panel shows the average return by quintiles of wealth, while the right panel shows the return pass-through. Both panels confirm the main results: Both the average return and the pass-through of returns are increasing in the level of wealth.

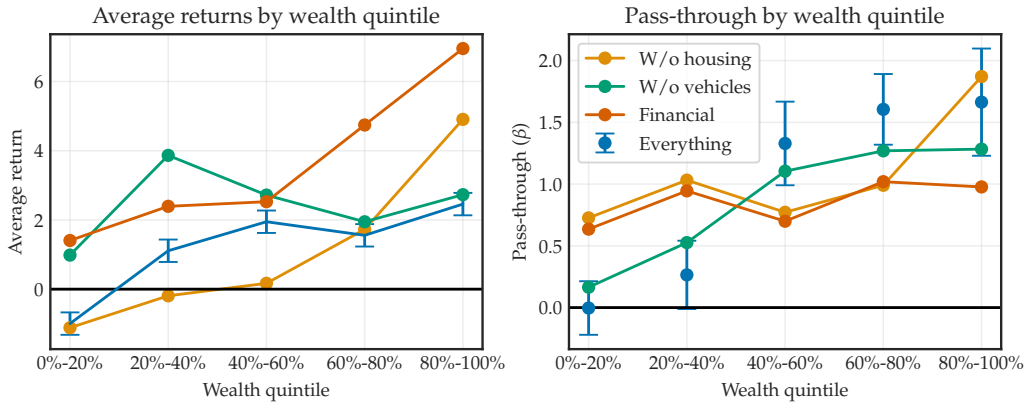


Figure A.5: Robustness of Returns Variables

Note: The figure shows robustness to different returns variables. The left panel shows the average return by quintiles of wealth. The right panel shows the return pass-through.

A.4.9 Portfolio and Normalized Returns

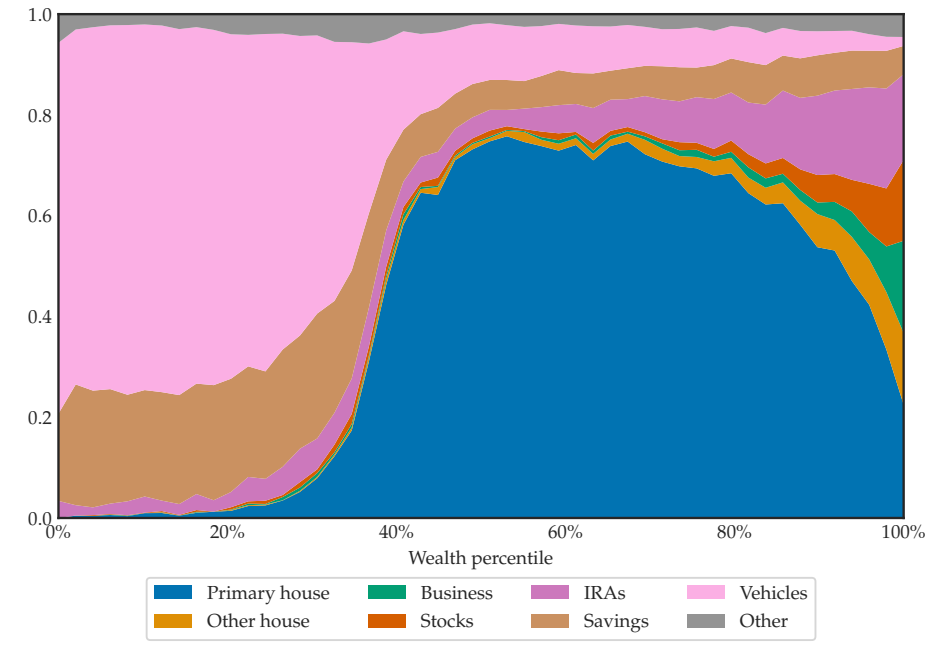


Figure A.6: Portfolio Shares Across Wealth Distribution

Note: The figure shows portfolio shares across the wealth distribution in the PSID.

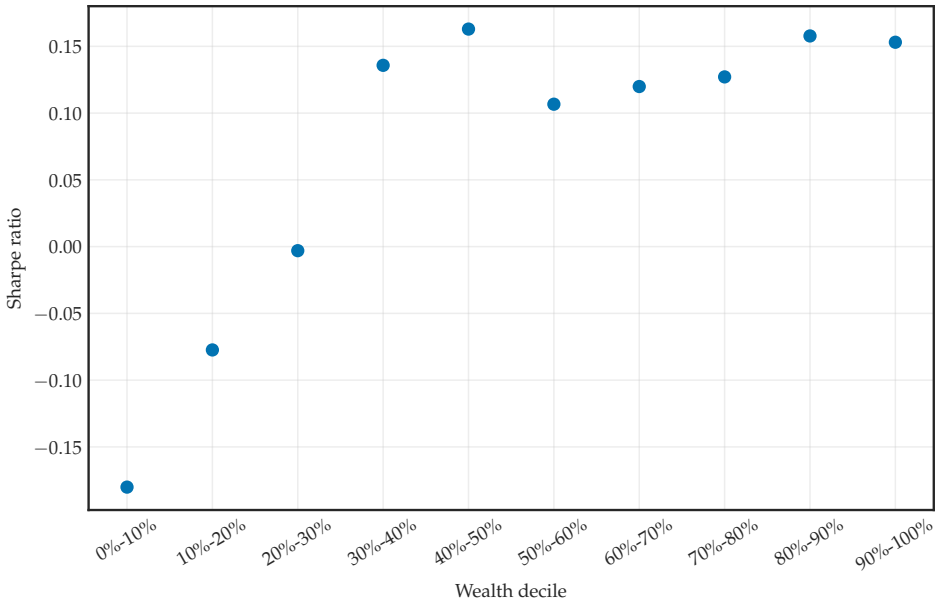


Figure A.7: Average Returns Normalized by Their Standard Deviation

Note: The figure shows the average returns normalized by their standard deviation across wealth deciles.

A.5 Returns by Wealth Quintiles in the SCF

I do not measure returns directly in the SCF, but I can still get at returns by wealth quintiles. I do so in the following way. Using the data provided by Kuhn and Rios-Rull (2016), I have portfolio shares by wealth quintiles. I then compute returns by wealth quintile by taking a weighted average of the returns of different asset classes provided by Jordà et al. (2019). The resulting time series are given in Figure A.8.

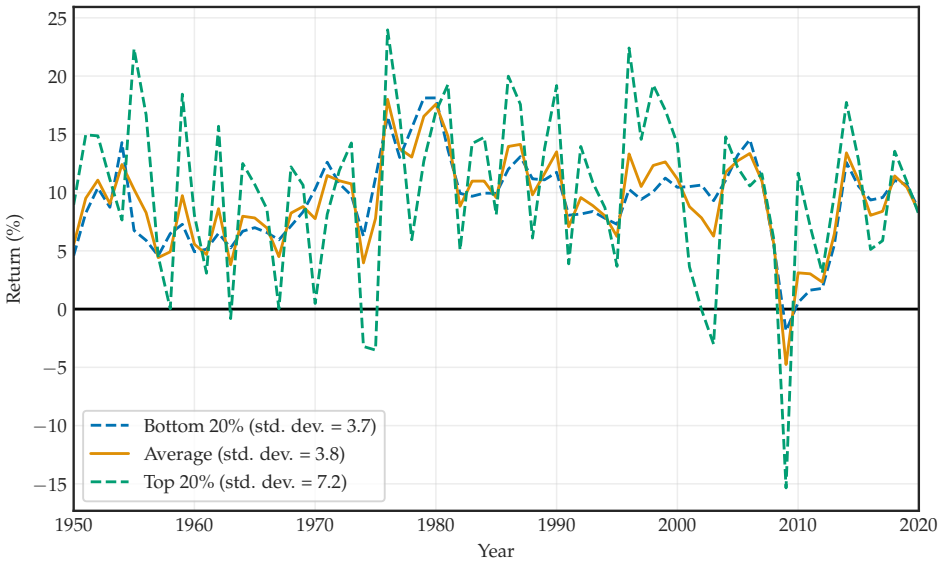


Figure A.8: Time Series of Average Returns in the SCF

Note: The figure shows a time series of average returns based on the SCF data.

While I do not have a panel of returns, I can still produce a figure similar to Figure A.9. In particular, I simply estimate the following time series regression,

$$r_t = \alpha^{(q)} + \beta^{(q)} r_t^{(q)} + \varepsilon_t^{(q)}, \tag{16}$$

where $r_t^{(q)}$ is the return for a particular quintile and r_t is the average return. I construct the average return as the average of the returns for the five quintiles. The resulting estimates are given in Figure A.9.

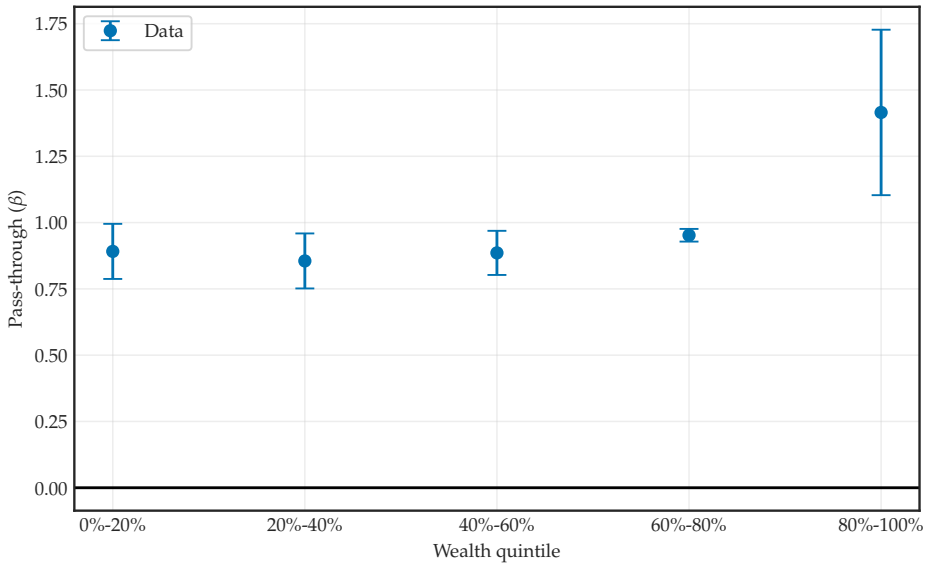


Figure A.9: Pass-Through of Average to Returns by Wealth in the SCF

Note: The figure shows estimates of β from eq. (16). The standard errors are robust to heteroskedasticity and autocorrelation (HAC) with 3 lags. Confidence intervals are at the 95% level.

B Appendix to Section 3

B.1 Solution Method and Equilibrium

I solve the households' problem using the endogenous gridpoint method (EGM) of Carroll (2006). I then use the methods from Auclert et al. (2021) to compute the Jacobians of the model. Lastly, I solve for the non-linear transition path under perfect foresight in the sequence space using a numerical equation solver (Boppart et al. 2018, Auclert et al. 2021).[30] This is equivalent to linearizing the model without perfect foresight (i.e., without aggregate risk) when shocks are small due to certainty equivalence. Note that while households have perfect foresight with respect to aggregate variables, they do not with respect to idiosyncratic variables, i.e., their labor income and returns.

Introducing heterogeneous returns changes the solution of the model compared to a standard HANK model in a few ways. First, it introduces a new state variable: The idiosyncratic return, e^r_{it}. Furthermore, heterogeneous returns stretch out the wealth distribution significantly. For this reason, it is important to have (i) a wealth grid with a very large maximum value, (ii) many grid points, and (iii) a highly non-linear grid. In particular, I (i) set the maximum wealth to 10^{14}, (ii) consider 401 grid points, and (iii) use the affine-exponential grid from Gouin-Bonenfant and Toda (2023). I have verified that the model has a stationary distribution, see Appendix C.2.

Despite these differences, the solution of the model is not much slower than the standard HANK model. This is because the model does not introduce any new choice variables. Additionally, the grid for returns does not have to have very many grid points.[31] Thus, I can still use the fast EGM method of Carroll (2006) to solve the household problem with heterogeneous returns.

I define the equilibrium of the model in Definition 1

Definition 1 (Equilibrium)**.** *Given sequences for* $\{\varepsilon_t, T_t, G_t\}$*, an initial household distribution,* $\mathcal{D}_0(a, e^z, e^r)$*, a competitive equilibrium is a path of household policies* $\{c_t(a, e^z, e^r), a_t(a, e^z, e^r)\}$*, distributions* $\mathcal{D}_t(a, e^z, e^r)$*, prices,*

$$\left\{P_t, W_t, w_t, \pi_t, \pi^W_t, i_t, r_t, p_t, r^a_t, \tilde{r}^a_t\right\},$$

30. The code is written in Python and based on the GEModelTools package.

31. I use 7 grid points for both the labor income process and the returns process, i.e., $7^2 = 49$ grid points for both in total.

and quantities,

$$\{B_t, \mathcal{T}_t, Y_t, N_t, D_t, Z_t, \tau_t, A_t, C_t\},$$

such that all households and firms optimize, monetary and fiscal policies follow their rules subject to constraints, and the goods market clears.

B.2 Recursive Formulation

Given sequences of aggregates, $(r_t^a)_{t=0}^{\infty}$, $(Z_t)_{t=0}^{\infty}$, $(T_t)_{t=0}^{\infty}$, and $(\tau_t)_{t=0}^{\infty}$, a household with state (a, e), where $e = (e^z, e^r)$, solves the following recursive problem:

$$\begin{aligned}
V_t(a,e) &= \max_{c,a'} u(c) + \beta\, \mathbb{E}_t\left[V_{t+1}(a', e')\right], \\
&\text{s.t.} \\
c + a' &= (1 + r^a)a + z + T_t - t, \\
z &= e^z Z + e^z \beta^z (Z_t - Z), \\
r^a &= r^a + e^r + \beta^r (r_t^a - r^a), \\
t &= \tau_t (r^a a + z), \\
e^z &\sim \text{Markov}(\mathcal{S}_z, \mathcal{P}_z), \\
e^r &\sim \text{Markov}(\mathcal{S}_r, \mathcal{P}_r).
\end{aligned}$$

I consider two modifications to do with changing the discount factor, β. The first is introducing permanent discount factor heterogeneity, which simply amounts to adding β as a permanent state.

B.3 Proof of Proposition 1

Since transfers do not change, the change in total income for household i at time $t = 0$ is given by

$$d\psi_{i0} = dx_{i0} + dz_{i0},$$

where $x_{it} \equiv r_{it}^a a_{it-1}$. The change in income going to this household is then

$$\frac{d\psi_{i0}}{d\Psi_0} = \frac{dx_{i0} + dz_{i0}}{d\Psi_0} = \frac{dx_{i0}}{d\Psi_0} + \frac{dz_{i0}}{d\Psi_0}.$$

Furthermore, assuming $dX_0 \neq 0$ and $dZ_0 \neq 0$:

$$\frac{d\psi_{i0}}{d\Psi_0} = \frac{dX_0}{d\Psi_0}\frac{dx_{i0}}{dX_0} + \frac{dZ_0}{d\Psi_0}\frac{dz_{i0}}{dZ_0}.$$

Write now $x_{i0} = r^a_{i0} a_{i,-1}$, such that $dx_{i0} = a_{i,-1} dr^a_{i0}$. Thus,

$$\frac{d\psi_{i0}}{d\Psi_0} = \frac{dX_0}{d\Psi_0}\frac{a_{i,-1}}{A_{-1}}\frac{dr^a_{i0}}{d\tilde{r}^a_{i0}} + \frac{dZ_0}{d\Psi_0}\frac{z_{i0}}{Z_0}\frac{\frac{dz_{i0}}{z_{i0}}}{\frac{dZ_0}{Z_0}}.$$

Using the definition of α_z and neglecting the $t = 0$ subscript yields the result.

B.4 The Elasticity of Earnings

I start by showing that the elasticity of z_{it} with respect to Z_t is β^z_{it}. To see this, note the definition of an elasticity:

$$\frac{\partial z_{it}}{\partial Z_t}\frac{Z_t}{z_{it}} = e^z_{it}\beta^z_{it}\frac{Z_t}{z_{it}}. \tag{17}$$

Evaluating this in the steady state yields that this elasticity is β^z_{it}.

Let me now discuss how to ground β^z_{it} empirically. Here, I turn to Guvenen et al. (2017). They estimate a regression that recovers something very similar to the elasticity, β^z_{it}. The only difference is that their elasticity is of earnings with respect to GDP and not aggregate earnings. However, these two elasticities are the same in the model. To see this, note that

$$Z_t = w_t N_t = \frac{W_t}{P_t} Y_t = \frac{1}{\mu} Y_t.$$

Thus, earnings are given by

$$z_{it} = \frac{e^z_{it}}{\mu} Y_{ss} + \frac{e^z_{it}}{\mu}\beta^z_{it}(Y_t - Y_{ss}).$$

Repeating the calculations in eq. (17) for Y_t instead of Z_t then yields that the elasticity of z_{it} with respect to Y_t is β^z_{it}. Thus, I use the estimates of β^z_{it} from Guvenen et al. (2017). In particular, I use their estimates of β^z_{it} as a function of earnings percentiles for males. To do so, I look at percentiles of the 7 grid points for e^z_{it} in the model and interpolate between the estimates in Guvenen et al. (2017). Finally, I divide all β^z_{it} by the same value such that $\int e^z_{it}\beta^z_{it}\, di = 1$, which ensures that $\int z_{it}\, di = Z_t$ for all t.

B.5 Robustness to Returns Process

I now consider three alternative specifications for the returns process. The first specification is that e^r_{it} follows a mean-zero AR(1), discretized by the Rouwenhorst method. I set the standard deviation and autocorrelation to match the baseline specification. This specification has three problems. First, it features a much too low wealth concentration, cf. Table A.4. Second, it features far too much negative income compared to the data, where there is essentially no negative income whatsoever. This negative income occurs because some very rich households get unlucky and earn negative returns. The baseline process almost entirely avoids this because households can only transition up or down one state. Third, it features too low scale dependence, measured as the difference between returns of the top and bottom 20%, cf. Table A.4.

The second specification is that there are 2 permanent returns, i.e., $e^r_{it} = e^r_i \in \{e^r_{\text{low}}, e^r_{\text{high}}\}$. The values are pinned down by (i) having a mean of zero, and (ii) matching the standard deviation of the baseline specification. This specification has two main problems. First, it features too little wealth concentration as measured by the top 1% wealth share, cf. Table A.4. Second—and in contrast with the AR(1)—it features too *much* scale dependence.

The third specification is a schedule for returns. In particular, I set the idiosyncratic returns process as a function of wealth: $e^r_{it} = f(a_{it-1})$. I choose this schedule to match the schedule in the baseline model. By construction—and in contrast with the other two specifications—it has a realistic degree of scale dependence. However, this comes at the cost of too low a standard deviation of returns.

In addition to these returns processes, I also consider one additional change: Correlated earnings and returns. This is motivated by the fact that there is empirical evidence that returns and earnings are correlated in the data, see Daminato and Pistaferri (2024). To add correlated returns, note that $\mathcal{P}_z$ is the transition matrix for e^z_{it}, while $\mathcal{P}_r$ is the transition matrix for returns e^r_{it}. The main specification obtains the transition matrix for the joint process under the assumption of independence, i.e., as

$$\mathcal{P}^{\text{ind}} = \mathcal{P}_r \otimes \mathcal{P}_z.$$

Instead of doing this, I now use a Gaussian copula (Nielsen 2006) governed by a single parameter, $\rho \in (-1, 1)$, to make the two correlated, yielding a different $\mathcal{P}^{\text{corr}}$. ρ controls the degree of correlation, nesting the independent case, $\mathcal{P}^{\text{corr}} = \mathcal{P}^{\text{ind}}$, when $\rho = 0$. Crucially, this method maintains the marginal distributions of both returns and earnings. I set $\rho = 0.5$, which yields a correlation coefficient of 0.21 between earnings

and returns. This specification performs reasonably, but has downsides compared to the baseline. In particular, both the wealth concentration at the top and the 1st income percentiles are slightly too small. More critically, the scale dependence—which is already slightly large compared to the data—is even larger with correlated returns.

In conclusion, the baseline specification provides the best fit to the data. However, let me emphasize that combining the different specifications would probably provide the best fit and most realistic specification at the cost of parsimony and without improving the fit much compared to the baseline specification.

	Baseline	AR(1)	Permanent	Schedule	Correlated
Average return	−1.3%	−1.3%	−1.3%	0.3%	−1.3%
Std. dev. of returns	9.3%	9.3%	9.3%	1.6%	9.3%
Top 1% wealth share	33.1%	15.7%	13.6%	14.6%	31.8%
Scale dependence	4.9%	−0.3%	9.4%	4.1%	6.2%
1st income percentile	0.03	−1.27	0.18	0.18	−0.03

Table A.4: Alternative Returns Processes

Note: The table shows descriptive statistics from models with alternative returns processes. "Scale dependence" refers to the difference in returns of the top 20% and bottom 20% in the wealth distribution.

B.6 Walras' Law

In this appendix, I show that the goods market clearing condition in eq. (14) implies asset market clearing. To start, note that aggregating households' budget constraints in eq. (4) yields

$$C_t + A_t = (1 + \tilde{r}_t^a)A_{t-1} + Z_t + T_t - \mathcal{T}_t.$$

Inserting the government's budget constraint in eq. (8) gives

$$C_t + A_t = (1 + \tilde{r}_t^a)A_{t-1} + Z_t + B_t - (1 + r_{t-1})B_{t-1} - G_t.$$

Using eq. (13) implies that

$$C_t + A_t = p_t + D_t + Z_t + B_t - G_t.$$

Using $D_t = Y_t - Z_t$ yields

$$C_t + A_t = p_t + Y_t + B_t - G_t.$$

Goods market clearing in eq. (14) then implies that

$$A_t = p_t + B_t,$$

which is exactly asset market clearing.

C Appendix to Section 4

C.1 Returns Process

The idiosyncratic part of returns, e^r_{it}, follows a discrete Markov chain with 7 grid points. These are set as the median within 7 bins, yielding the grid points for e^r_{it}. The ergodic distribution over these grid points, π, is taken as the empirical distribution.

What remains to be specified is the transition matrix. I parametrize the transition matrix as follows. For each of the 7 states, there is some probability of staying in that state. Additionally, there is some probability of going up one state and the same probability of going down one state. All other transitions have zero probability. If I assumed that it was possible to jump multiple steps, some ultra-rich households would suddenly earn a negative return. This would mean that they get (very) negative income. This is at odds with the data, where essentially no households earn negative income. With this specification, the transition matrix can be written as

$$\mathscr{P}_r = \begin{pmatrix} p_1 & 1-p_1 & 0 & 0 & 0 & 0 & 0 \\ \frac{1-p_2}{2} & p_2 & \frac{1-p_2}{2} & 0 & 0 & 0 & 0 \\ 0 & \frac{1-p_3}{2} & p_3 & \frac{1-p_3}{2} & 0 & 0 & 0 \\ 0 & 0 & \frac{1-p_4}{2} & p_4 & \frac{1-p_4}{2} & 0 & 0 \\ 0 & 0 & 0 & \frac{1-p_5}{2} & p_5 & \frac{1-p_5}{2} & 0 \\ 0 & 0 & 0 & 0 & \frac{1-p_6}{2} & p_6 & \frac{1-p_6}{2} \\ 0 & 0 & 0 & 0 & 0 & 1-p_7 & p_7 \end{pmatrix}.$$

There is the following relationship between the transition matrix and the stationary distribution,

$$p_i = \begin{cases} 1 - \frac{c}{\pi_i} & \text{for} \quad i = 1,7 \\ 1 - \frac{2c}{\pi_i} & \text{for} \quad i = 2,3,4,5,6 \end{cases},$$

for some parameter c satisfying

$$0 < c \leq \min\left(\pi_1, \frac{\pi_2}{2}, \ldots, \frac{\pi_6}{2}, \pi_7\right).$$

Thus, a choice of c and taking π from the data pins down the full transition matrix, $\mathscr{P}_r$. c is closely related to the persistence of the process: As $c \to 0$, states are permanent. For higher values of c, the probability of changing state is higher.

In Figure A.10, I consider varying c and reporting different statistics of the model.

In particular, the figure shows absolute deviations of statistics in the model from their data equivalent. As the figure shows, there is a trade-off between increasing and decreasing c from its calibrated level in terms of the fit of the model: Increasing c improves its fit to the income and wealth distribution Gini coefficients slightly, while it hurts its fit to the top 0.1% share.

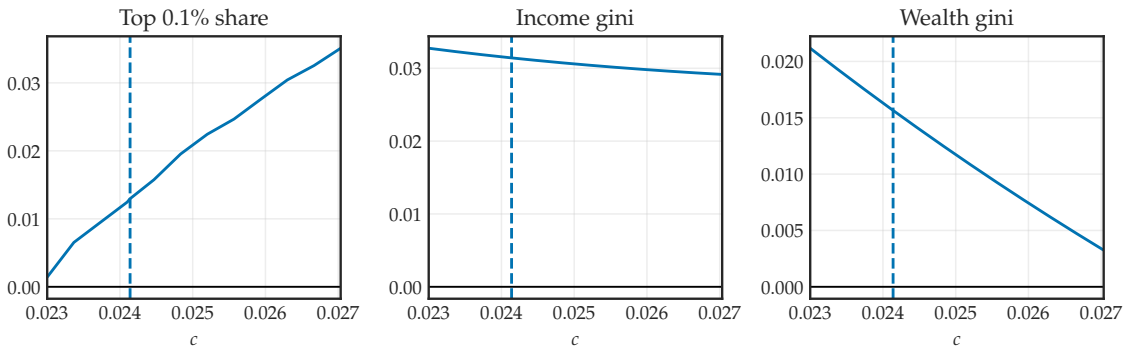


Figure A.10: Robustness of c

Note: The figure shows the absolute deviation of statistics in the model from their data counterparts as a function of c.

In Figure A.11, I consider measuring the riskiness of returns and their Sharpe ratio. The left panel shows the conditional standard deviation of returns, $\sqrt{\text{Var}(e^r_{it} \mid e^r_{it-1})}$, by quintiles of wealth. This answers the question: Given what my return is this year, what is the standard deviation of my returns next year? The second panel shows the unconditional cross-sectional standard deviation of returns, $\sqrt{\text{Var}(e^r_{it})}$, by quintiles of wealth. This answers the question: What is the standard deviation of returns? Both measures of riskiness of returns are increasing in wealth. Thus, richer households have riskier returns. The third panel computes the Sharpe ratio in a way equivalent to what I compute in the data, i.e., by dividing the average return by the cross-sectional standard deviation of returns (the second panel). This is increasing in wealth. All panels match the PSID data qualitatively. Crucially, this is untargeted.

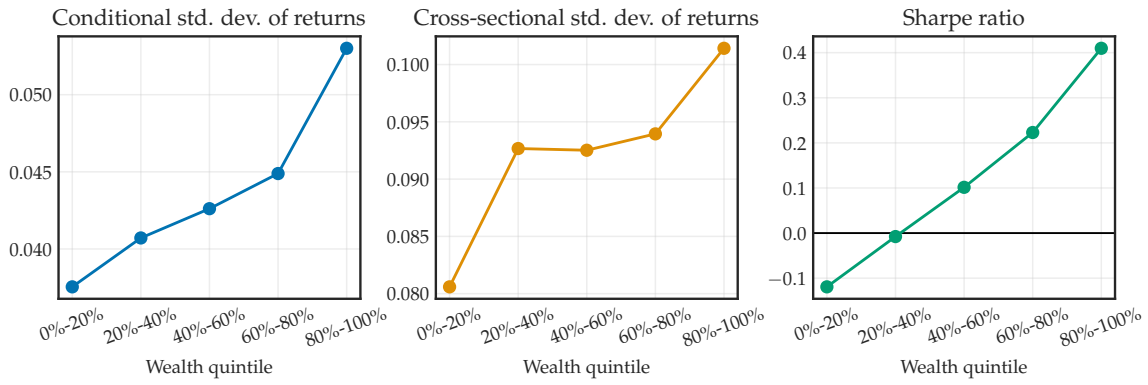


Figure A.11: Riskiness of Returns and Their Sharpe Ratios

Note: The figure shows (i) the conditional standard deviation of returns, (ii) the unconditional cross-sectional standard deviation of returns, and (iii) the Sharpe ratios of returns. All three are computed by quintiles of wealth.

C.2 Stationarity

When adding heterogeneous returns to the otherwise standard incomplete markets model, it is no longer trivial whether the distribution of households is stationary. For this reason, I have verified that my calibration of the model is stationary. I have done so using four different approaches, presented here.

1. I have iterated on the distribution using the histogram method from various different initial distributions and verified that the distribution converges to the same one for all starting distributions.

2. I have solved the model for different (i) maximum grid points in the asset grid and (ii) number of grid points. In all cases, I have verified that the resulting distributions are numerically indistinguishable. In particular, there exists some asset level (around $a_{it} = 10^9$ independent of the asset grid) for which there is (numerically indistinguishable from) zero probability mass above this level.

3. Benhabib et al. (2015) show that under certain conditions that ensure a stationary wealth distribution, the wealth distribution asymptotically has a Pareto tail. I have verified empirically that my model also has a Pareto tail, which points to a stationary wealth distribution.

4. I have simulated panels of households and verified that the mean wealth always converges to and then fluctuates around the expectation of the stationary distribution of wealth.

C.3 Returns Pass-Through Functional Form Robustness

In this appendix, I consider different functional forms for the pass-through of average to household returns, β^r_{it}. The baseline specification is

$$\beta^r_{it} = \frac{\log\left\{1 + (1+\theta_0)(1-\theta_1^{a_{it-1}})^2\right\}}{\log(2+\theta_0)}.$$

I now consider the following four alternative specifications:

$$\begin{aligned}
\beta^r_{it} &= \exp\left\{\theta_0\left[1 - \exp(-\theta_1 a_{it-1})\right]\right\}, \\
\beta^r_{it} &= \frac{2}{1+\exp(\theta_0 - \theta_1 a_{it-1})}, \\
\beta^r_{it} &= \frac{(1-\theta_1^{a_{it-1}})^2}{1+\theta_0(1-\theta_1^{a_{it-1}})^2}, \\
\beta^r_{it} &= \exp\left\{\frac{\theta_0 a_{it-1}}{\sqrt{1+(\theta_1 a_{it-1})^2}}\right\}.
\end{aligned}$$

where all are normalized to average to 1 in steady state. I estimate all five functional forms and report the fit in Figure A.12.

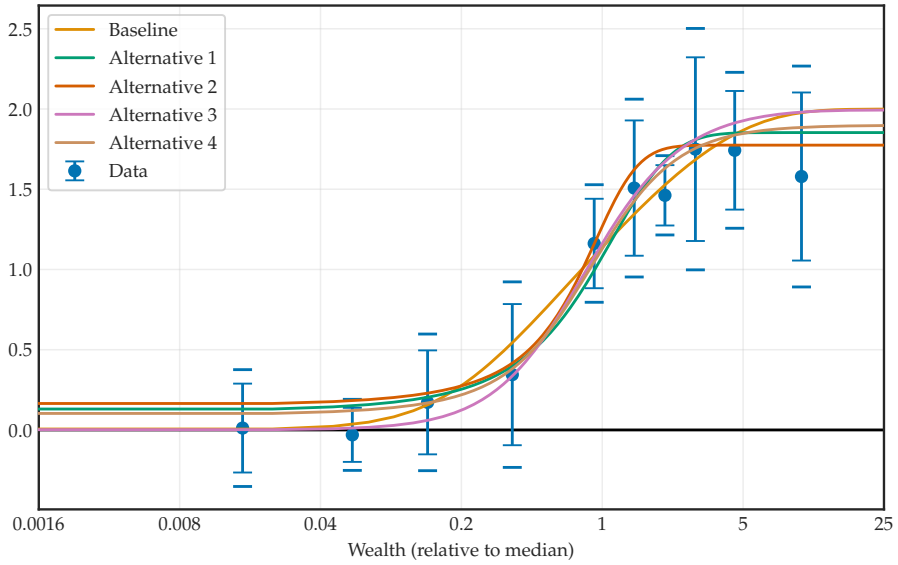


Figure A.12: Robustness of Pass-Through of Average to Households' Returns by Wealth

Note: The figure shows estimates of β from eq. (3) by wealth, a_{it-1}, for five different functional forms. The standard errors are clustered by year. Ticks indicate 95% and 99% confidence intervals.

I report the impulse response of output, the top 0.1% income share due to monetary

policy, and the objective function in Table A.5. I note that the baseline specification has the best fit, which is why I chose it. However, all other functional forms have good fits as well. Crucially, they all imply almost exactly identical responses, so my results are very robust to the functional form of β_{it}^r.

Model	Output increase	Top 0.1% income share	Objective
Baseline	1.450%	10.67%	0.00013
Alternative 1	1.467%	10.62%	0.00033
Alternative 2	1.480%	10.59%	0.00043
Alternative 3	1.455%	10.65%	0.00022
Alternative 4	1.460%	10.64%	0.00022

Table A.5: Robustness

Note: The table shows the robustness of the results to changes in the functional form of β_{it}^r.

D Appendix to Section 5

D.1 Scale Dependence in the Model and the Data

As shown in Figure 9, returns are increasing in the level of wealth in both the model with heterogeneous returns and the PSID data. This is untargeted in the model and occurs endogenously. However, Figure 9 also shows that the returns in the model are somewhat more increasing in wealth than in the data. There are three reasons I consider this reasonable. First, measurement error in returns will make the relationship look weaker in the data than it actually is. Second, if the relationship between returns and wealth was weaker in the model, the fit to Figure 10 would be worse. Third, the relationship between returns and wealth I find is weaker than most other papers, as I show in Figure A.13.

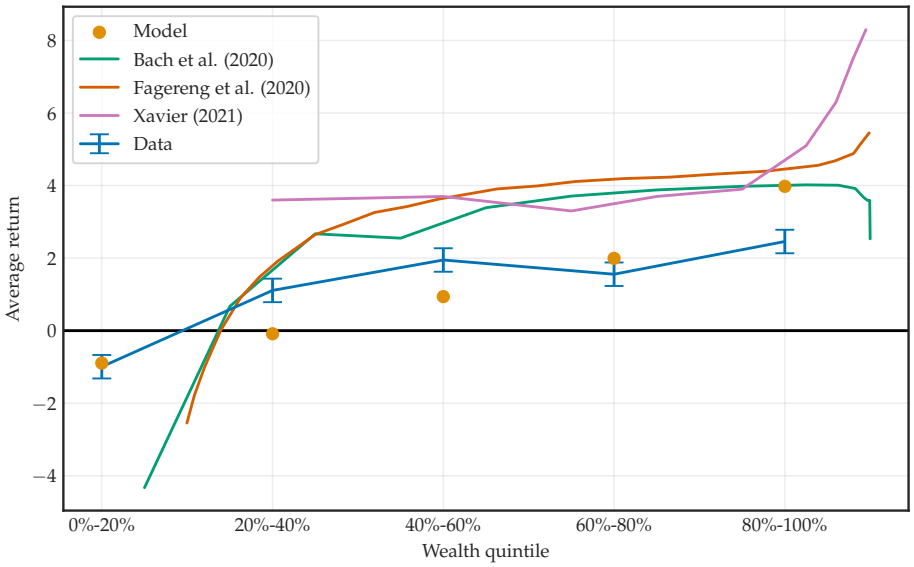


Figure A.13: Average Return by Wealth (Scale Dependence) In Model and Different Data

Note: The figure shows the average return by quintiles of a_{it-1}. The papers are Bach et al. (2020), Fagereng et al. (2020), and Xavier (2021).

D.2 Household Consumption Policy Functions

In this Appendix, I show how the policy functions for households depend on their return. To be specific, the policy function for consumption in the ergodic steady state

can be written as[32]

$$c_{it} = c(a_{it-1}, e^z_{it}, e^r_{it}).$$

Instead of plotting directly the policy functions, I plot a more easily interpretable object: The MPC. I define the MPC as the marginal increase in consumption from a marginal increase in cash-on-hand:

$$\text{mpc}_{it}(a_{it-1}, e^z_{it}, e^r_{it}) = \frac{\partial c(a_{it-1}, e^z_{it}, e^r_{it})}{\partial a_{it-1}} \frac{1}{1 + (1 - \tau_t) r^a_{it}}.$$

The factor $(1 + (1 - \tau_t) r^a_{it})^{-1}$ simply adjusts for the factor that a unit increase in wealth increases cash-on-hand not by a unit, but by $1 + (1 - \tau_t) r^a_{it}$. This ensures that the MPC is between 0 and 1.

I plot the MPC policy function in Figure A.14. In particular, I fix a value of e^z_{it} and then plot consumption as a function of wealth for the 7 different values of e^r_{it} and the corresponding returns. The figure shows that households with higher returns have a much lower MPC, instead saving more of income windfalls. This is what creates scale dependence.

32. The subscript t refers to variables for households changing even in the ergodic steady state. Only aggregate variables are fixed at their steady state values. This also explains why there is no subscript t on $c(\cdot)$.

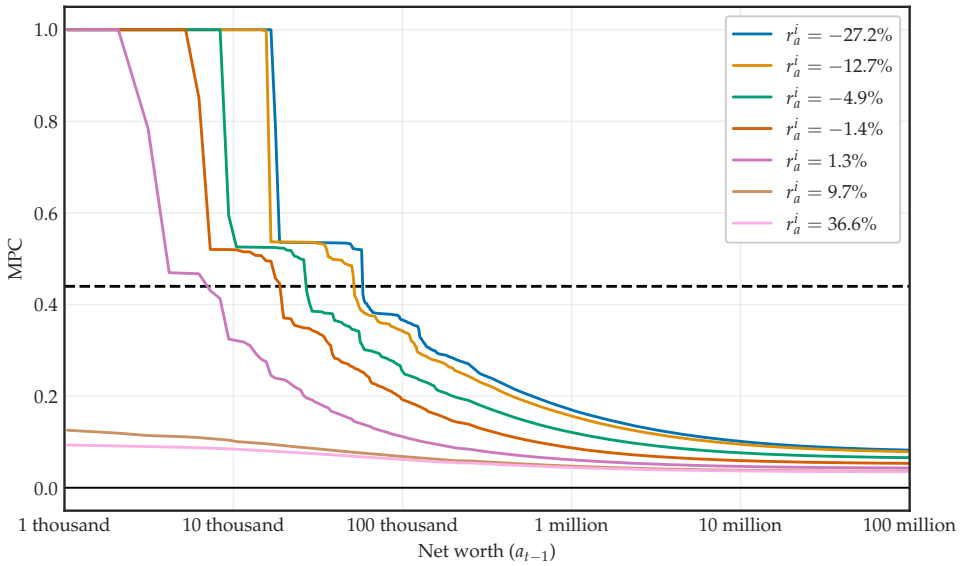


Figure A.14: MPC Policy Functions

Note: The figure shows the MPC policy function in the ergodic steady state for the model with heterogeneous returns. In particular, I fix a value of e_{it}^z and plot consumption as a function of wealth for different values of returns.

D.3 Lorenz Curves

Figures A.15–A.16 show the Lorenz curves for wealth, income, and earnings in the model and the data.

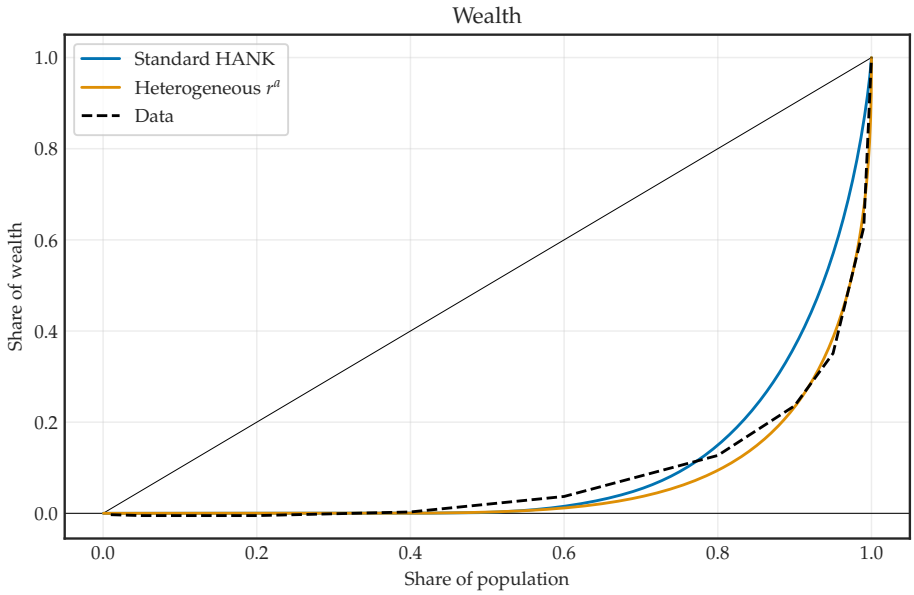


Figure A.15: Lorenz Curves for Wealth

Note: The figure shows the Lorenz curves of wealth in both models and the 2019 SCF.

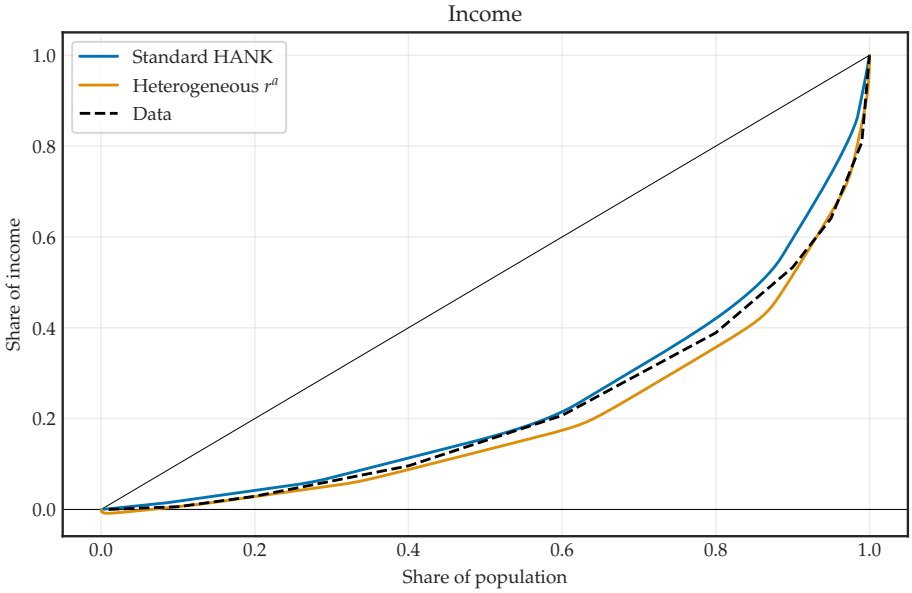


Figure A.16: Lorenz Curves for Income

Note: The figure shows the Lorenz curves of income in both models and the 2019 SCF.

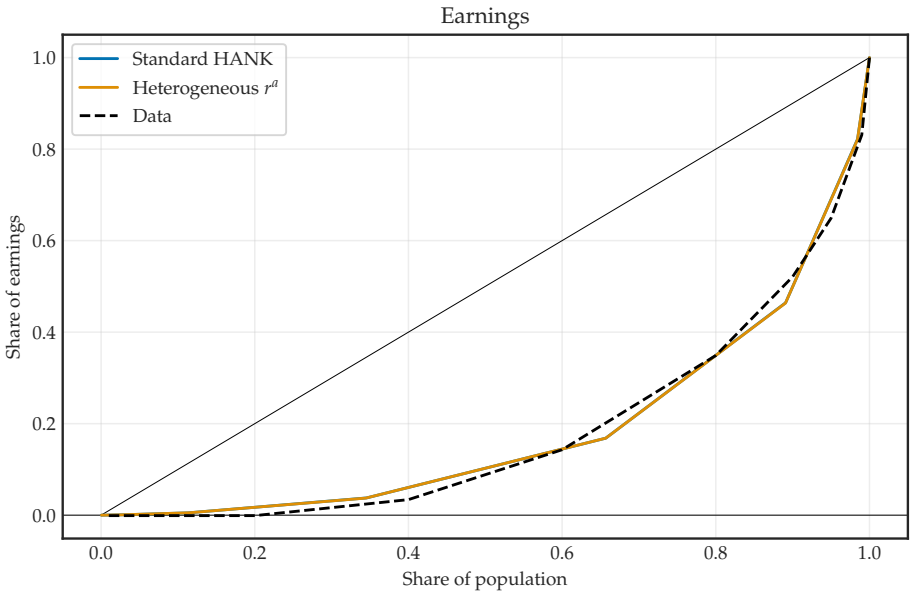


Figure A.17: Lorenz Curves for Earnings

Note: The figure shows the Lorenz curves of earnings in both models and the 2019 SCF.

D.4 The Concentration of Household Variables

One feature of the distributions of household variables in the data is that they follow an ordering of concentration. In particular:

$$g(c_{it}) < g(a_{it}) < g(x_{it}), \tag{18}$$

where $g(\cdot)$ denotes a measure of concentration (inequality) like the Gini index or the Pareto tail index. Gaillard et al. (2023) study this in a standard heterogeneous agents model with common returns and find that the tail index of all three variables is the same, in contrast with the data.[33]

How does my model perform regarding the ranking in (18)? To study this, I report top shares of the three variables in Table A.6. I find that my model satisfies the ranking: Consumption is the most equal, capital income is the least equal, and wealth is somewhere in between.

Variable	Top 5%	Top 1%	Top 0.1%	Top 0.01%
Capital income	107%	71%	30%	11%
Wealth	61%	33%	13%	5%
Consumption	25%	10%	3%	1%

Table A.6: The Concentration of Household Variables

Note: The table shows top shares of selected household variables in the model with heterogeneous returns.

33. I do not consider earnings here, as it has a simple 7-point discrete distribution implied by the Markov chain.

E Appendix to Section 6

E.1 Additional IRFs to Monetary Policy

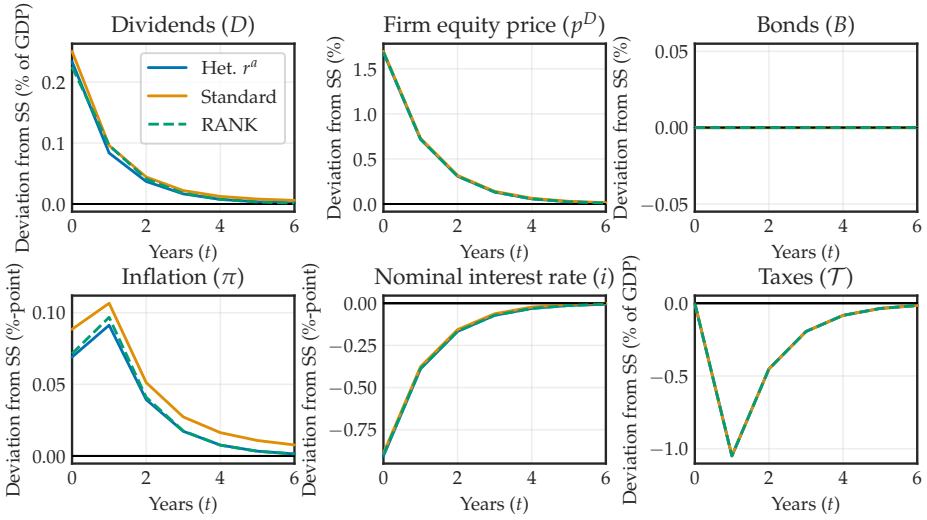


Figure A.18: Additional Aggregate Effects of Monetary Policy

Note: See Figure 13.

E.2 Understanding the Transmission of Monetary Policy

The fact that aggregate outcomes are different can hide a change in the *transmission* of monetary policy. To see this, I now decompose consumption into the channels through which monetary policy works. This decomposition is similar to the one in Kaplan et al. (2018). The decomposition is in Proposition 2.

Proposition 2. *Consider a monetary policy shock. The response of consumption is given by*

$$dC = \underbrace{\boldsymbol{M}^{r} d\boldsymbol{r}}_{\text{1. Direct}} + \underbrace{\boldsymbol{M}^{\boldsymbol{Z}} d\boldsymbol{Z}}_{\text{2. Labor}} + \underbrace{\boldsymbol{M}^{\boldsymbol{\tau}} d\boldsymbol{\tau}}_{\text{3. Taxes}} + \underbrace{\boldsymbol{M}^{\boldsymbol{X}} dX_0}_{\text{4. Revaluation}} + \underbrace{\boldsymbol{M}^{Cov} dCov\left(r_i^a, \frac{a_{i,-1}}{A_{-1}}\right)}_{\text{5. Redistribution}},$$

where the $\boldsymbol{M}$'s are the Jacobians defined below and $X_0 = \int r_{i0}^a a_{i,-1}\, di$.

Proof. The sequence space consumption function can be written as

$$C_t = C_t(\{Z_s\}_{s=0}^{\infty}, \{r_s^a\}_{s=0}^{\infty}, \{\tau_s\}_{s=0}^{\infty}, \{T_s\}_{s=0}^{\infty}).$$

Linearizing this in the sequence space gives

$$d\boldsymbol{C} = \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{Z}} d\boldsymbol{Z} + \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{r^a}} d\boldsymbol{r^a} + \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{\tau}} d\boldsymbol{\tau},$$

ignoring changes in lump-sum transfers, T_t, as I am considering a monetary policy shock. The sequence Jacobians of this function are then the partial derivative of the consumption function at time t with respect to one of the inputs at time s. These are presented by the matrices $\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{Z}}$, $\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{r^a}}$, and $\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{\tau}}$. Figure A.19 shows selected columns of the sequence space Jacobians for Z_s (the i-MPC matrix) and r_s^a for the model with heterogeneous returns and the standard HANK model.

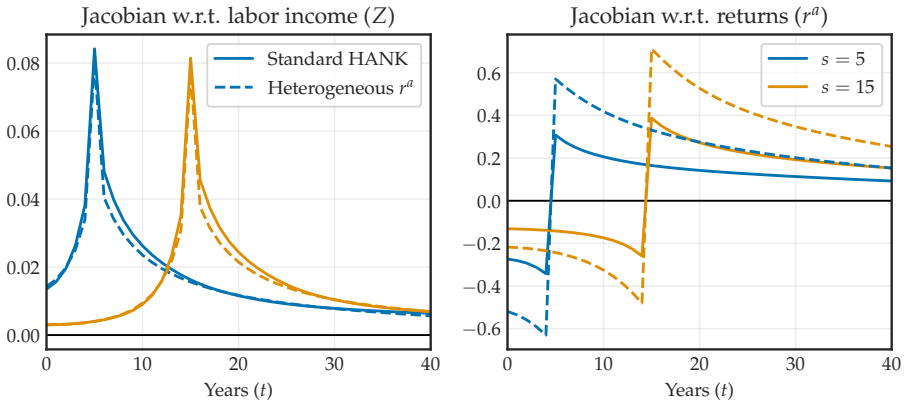


Figure A.19: Sequence Space Jacobians of Household Problem

Note: The figure shows two columns ($s = 5$ and $s = 15$) of the sequence space Jacobians of the household problem for real labor income and returns, i.e., $\partial C_t / \partial Z_s$ and $\partial C_t / \partial r_t^a$. It does so for both the model with heterogeneous returns and the standard HANK model.

To proceed with the proof, note that

$$d\tilde{\boldsymbol{r}}^{\boldsymbol{a}} = \boldsymbol{L} d\boldsymbol{r} + \iota d\tilde{r}_0^a,$$

where $\boldsymbol{L}$ is the lag operator and $\iota = (1, 0, 0, \dots)'$. Note also that

$$d\tilde{r}_t^a = dr_t^a + d\mathrm{Cov}\left(r_{it}^a, \frac{a_{it-1}}{A_{t-1}}\right).$$

Combining these two expressions yields

$$d\boldsymbol{r}^{\boldsymbol{a}} = \boldsymbol{L} d\boldsymbol{r} + \iota d\tilde{r}_0^a - d\mathrm{Cov}\left(r_i^a, \frac{a_{i,-1}}{A_{-1}}\right),$$

Inserting this into the consumption function yields

$$d\boldsymbol{C} = \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{r}^a}\boldsymbol{L}d\boldsymbol{r} + \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{Z}}d\boldsymbol{Z} + \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{\tau}}d\boldsymbol{\tau} + \frac{1}{A_{ss}}\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{r}^a}\iota dX_0 - \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{r}^a}d\text{Cov}\left(r^a_{it}, \frac{a_{it-1}}{A_{t-1}}\right).$$

Rewriting slightly and defining Jacobians, I arrive at the equation:

$$d\boldsymbol{C} = \underbrace{\boldsymbol{M}^r d\boldsymbol{r}}_{\text{1. Direct}} + \underbrace{\boldsymbol{M}^{\boldsymbol{Z}} d\boldsymbol{Z}}_{\text{2. Labor}} + \underbrace{\boldsymbol{M}^{\boldsymbol{\tau}} d\boldsymbol{\tau}}_{\text{3. Taxes}} + \underbrace{\boldsymbol{M}^X dX_0}_{\text{4. Revaluation}} + \underbrace{\boldsymbol{M}^{\text{Cov}} d\text{Cov}\left(r^a_i, \frac{a_{i,-1}}{A_{-1}}\right)}_{\text{5. Redistribution}},$$

where

$$\begin{aligned}
\boldsymbol{M}^r &\equiv \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{r}^a}\boldsymbol{L}, \\
\boldsymbol{M}^{\boldsymbol{Z}} &\equiv \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{Z}}, \\
\boldsymbol{M}^{\boldsymbol{\tau}} &\equiv \frac{\partial \boldsymbol{C}}{\partial \boldsymbol{\tau}}, \\
\boldsymbol{M}^X &\equiv \frac{1}{A_{ss}}\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{r}^a}\iota, \\
\boldsymbol{M}^{\text{Cov}} &\equiv -\frac{\partial \boldsymbol{C}}{\partial \boldsymbol{r}^a}.
\end{aligned}$$

□

I decompose the response of consumption, dC_0, into these channels in the two models in Table A.7. Consider first the standard model with common returns. In this model, less than half of the response of consumption is driven by direct intertemporal substitution (channel 1). This channel reflects household pushing consumption forward in time and is the main channel operating in standard representative agent models. The remaining effect on consumption is through indirect channels: Higher labor income (channel 2), lower tax rates (channel 3), and higher capital income due to a revaluation of wealth (channel 4). Kaplan et al. (2018) make the point that these indirect channels are strong drivers of consumption in HANK models, which is also the case here.

	Standard HANK	Heterogeneous r^a
1. Direct	0.73	1.35
2. Labor	0.18	0.15
3. Taxes	0.47	0.43
4. Revaluation	0.20	0.42
5. Redistribution	0.00	−0.88
Total	1.57	1.48

Table A.7: Decomposition of Consumption in Response to Monetary Policy

Note: The table shows a decomposition of the response of consumption on impact to the monetary policy shock in two models: The standard HANK with common returns and the baseline HANK with heterogeneous returns.

Consider then the model with heterogeneous returns. Here, a new channel is active: The redistribution channel (channel 5). The redistribution channel lowers consumption. This happens because average returns change less. The left panel of Figure A.20 shows this, plotting the average return. This also explains why the direct and revaluation channels are stronger in Table A.7. For instance, the direct effect says what would happen to consumption due to changes in the real interest rate if there were no redistribution. In this counterfactual case, the average return would change as much as with common returns, which implies a stronger effect on consumption. Why does the average return change less? Because the shock redistributes to households that are wealthier, so the average return cannot change as much. The right panel of Figure A.20 shows this, reporting the change in returns for the rich and the poor. I note that this is exactly consistent with the data in Figure 5.

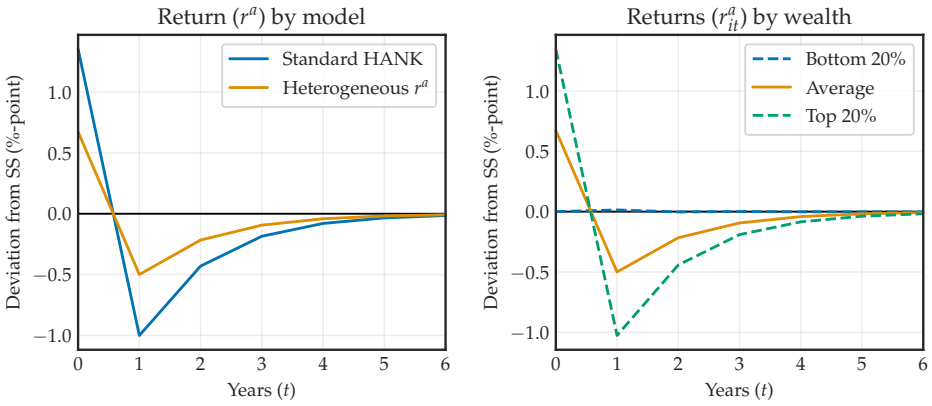


Figure A.20: The Effects of Monetary Policy on Returns Across Models and Wealth

Note: The left panel shows the impulse response function of the average return. The right panel shows impulse response functions of returns for different wealth quintiles. The x-axis shows years after the shock.

Let me emphasize that the change in transmission with heterogeneous returns is entirely due to β^r_{it}. To do so, I first consider the heterogeneous returns model with common pass-through, i.e. $\beta^r_{it} = 1$, but still heterogeneous returns. Second, on the other hand, I also consider a model where β^r_{it} varies more strongly over the wealth distribution. I do this by setting $\theta = (0, 0.5)$. This implies that β^r_{it} is ≈ 0 for poor households and much larger than the baseline specification for rich households. In this sense, β^r_{it} varies more strongly as a function of wealth. Figure A.21 shows this.

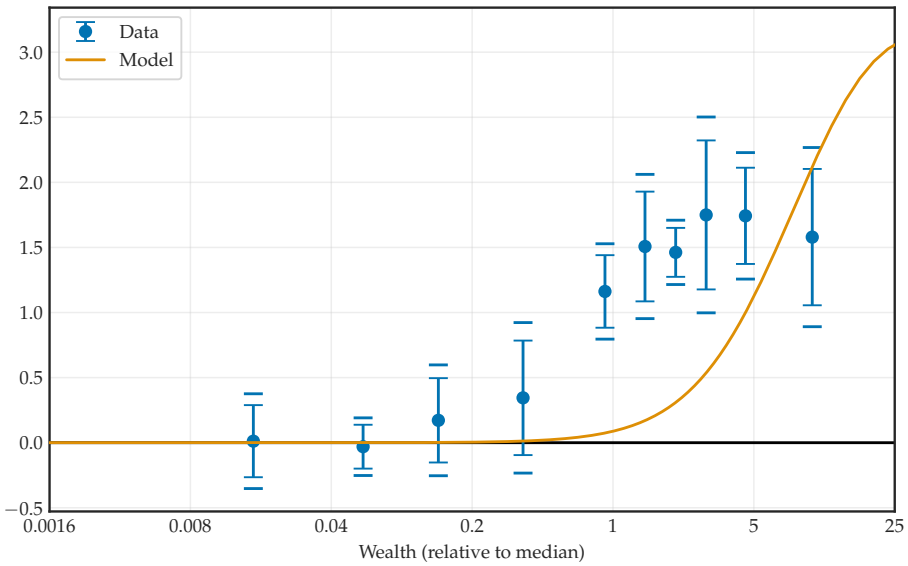


Figure A.21: More Varied Pass-Through of Aggregate to Households' Returns, β^r_{it}

Note: The x-axis shows the lagged asset distribution, while the y-axis shows β^r_{it} with $\theta = (0, 0.5)$ for households at these points in the wealth distribution. See Figure 7. Ticks indicate 95% and 99% confidence intervals.

Table A.8 shows the decomposition. The table shows a clear picture. First, with $\beta^r_{it} = 1$, the decomposition is essentially identical to a model with common returns. Second, with a more varied β^r_{it}, the redistribution and direct channels—which take opposite signs—are even stronger. On net, the redistribution channel overpowers, so consumption is smaller than in the baseline model. This shows that β^r_{it} *can* change the aggregate effects of monetary policy, though the difference between β^r_{it} for the rich and poor empirically is not strong enough to do this in any significant way.

	Common pass-through	More varied β^r
1. Direct	0.76	1.63
2. Labor	0.16	0.17
3. Taxes	0.43	0.47
4. Revaluation	0.21	0.65
5. Redistribution	0.00	−1.52
Total	1.56	1.39

Table A.8: Decomposition of Consumption in Response to Monetary Policy With Different Pass-Through

Note: The table shows a decomposition of the response of consumption on impact to the monetary policy shock in two models: A model with $\beta^r_{it} = 1$ and a model with $\theta = (0, 0.5)$.

The takeaway is that the decomposition of the aggregate effects of monetary policy with heterogeneous returns but common *pass-through*, i.e. $\beta^r_{it} = 1$, is almost identical to the decomposition with common returns. On the other hand, when the variation in $\beta^r_{it} = 1$ is made more drastic over the wealth distribution, the redistribution channel is stronger, dragging consumption down. This case is only theoretically interesting and not empirically relevant, so I do not pursue it further.

E.3 Relation to Werning (2015)

Werning (2015) shows that the response of aggregate consumption to real interest rate changes is identical in incomplete markets and complete markets models under certain assumptions. The models in this paper have the same neutrality result approximately, similar to Kaplan et al. (2018). However, the models in this paper

(and Kaplan et al. 2018) do not satisfy this assumption exactly—even with common returns.

I now consider simplifying my model such that it fits the setup in Werning (2015). In particular, I start with the fully calibrated model with heterogeneous returns. I then remove the government: $\tau_{ss} = B_{ss} = T_{ss} = 0$, which also implies $G_{ss} = 0$. I also remove the heterogeneous pass-through: $\beta^r_{it} = \beta^z_{it} = 1$. With this parametrization, I recalibrate β and r_{ss}. This gives me a simplified version of the model with heterogeneous returns. I also consider a version of this simplified model without heterogeneous returns by setting $e^r_{it} = 0$ and recalibrating β. Additionally, I consider a representative agent (RANK) version of the model, where consumption is described by the standard Euler equation:

$$u'(C_t) = \beta(1 + r_t)u'(C_{t+1}).$$

Figure A.22 shows the response of consumption in all three models. Note first that the response of consumption is completely identical in the standard HANK and RANK models. This is exactly Werning (2015). Note additionally that consumption in the heterogeneous returns model is *almost* the same, but not quite. The difference is small enough not to be economically meaningful but large enough not to be an artifact of the numerical solution.

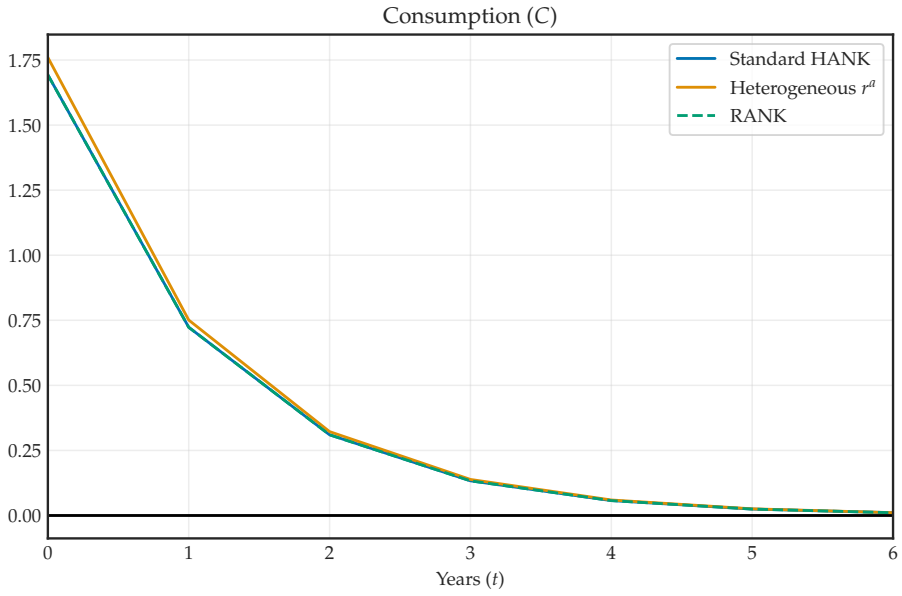


Figure A.22: IRFs to a Monetary Policy Shock in Simplified Models

Note: The figure shows impulse response functions to a 1 percentage point drop in the real interest rate in simplified versions of three models. The x-axis shows years after the shock.

To see the math, let $\tilde{c}_{it}$ denote the consumption of household i at time t in the absence of the monetary policy shock, i.e., when aggregate variables are in steady state. Werning (2015) shows that

$$\frac{c_{it}}{\tilde{c}_{it}} = \frac{C_t^{\text{RA}}}{C_{ss}^{\text{RA}}},$$

where C_t^{RA} is consumption in the RANK model faced with the same monetary policy shock. This does not hold with heterogeneous returns. Instead, it holds that

$$\frac{c_{it}}{\tilde{c}_{it}} = s_{it}\frac{C_t^{\text{RA}}}{C_{ss}^{\text{RA}}},$$

where s_{it} is simply defined as the scaling factor such that this holds. Aggregate consumption is then

$$C_t = \int c_{it}\, di = \int s_{it}\frac{C_t^{\text{RA}}}{C_{ss}^{\text{RA}}}\tilde{c}_{it}\, di = C_t^{\text{RA}}\int s_{it}\tilde{c}_{it}\, di = C_t^{\text{RA}}\left[S_t + \text{Cov}(s_{it}, \tilde{c}_{it})\right],$$

where $S_t = \int s_{it}\, di$. Under which assumptions is $C_t = C_t^{\text{RA}}$? The assumptions in Werning (2015) imply that $s_{it} = 1$ for all i and t. However, while $s_{it} = 1$ is *sufficient*, it is not *necessary*. For instance, one obtains $C_t = C_t^{\text{RA}}$ despite $s_{it} \neq 1$ if both (i) $S_t = 1$, and (ii) $\text{Cov}(s_{it}, \tilde{c}_{it}) = 0$. This is (approximately) the case in the model with heterogeneous returns. Intuitively, the consumption of households *is* affected differently by the monetary policy shock, but the way it is affected is (almost) independent of initial consumption, so all the changes wash out in the aggregate.

Figure A.23 illustrates this. The first row of the figure shows the change of consumption across households for both the standard HANK model and the model with heterogeneous returns. Each dot corresponds to one of the discretized states in the model. In the standard HANK model, each household's consumption rises in the same proportion, i.e., $s_{it} = 1$ as in Werning (2015). In the model with heterogeneous returns, there is a difference in how much consumption increases, with consumption increases ranging from around 1.2% to 2.4%. This clearly shows that the Werning (2015) result does *not* apply in the model with heterogeneous returns. However, the figure also shows that the changes in consumption are (almost) uncorrelated with the initial level of consumption. The second row makes this point, dividing the distribution of consumption into 20 quantiles of 5%. In this case, consumption increases by the same proportion for households at different points in the consumption

distribution. In other words, $\text{Cov}(s_{it}, \tilde{c}_{it}) \approx 0$.

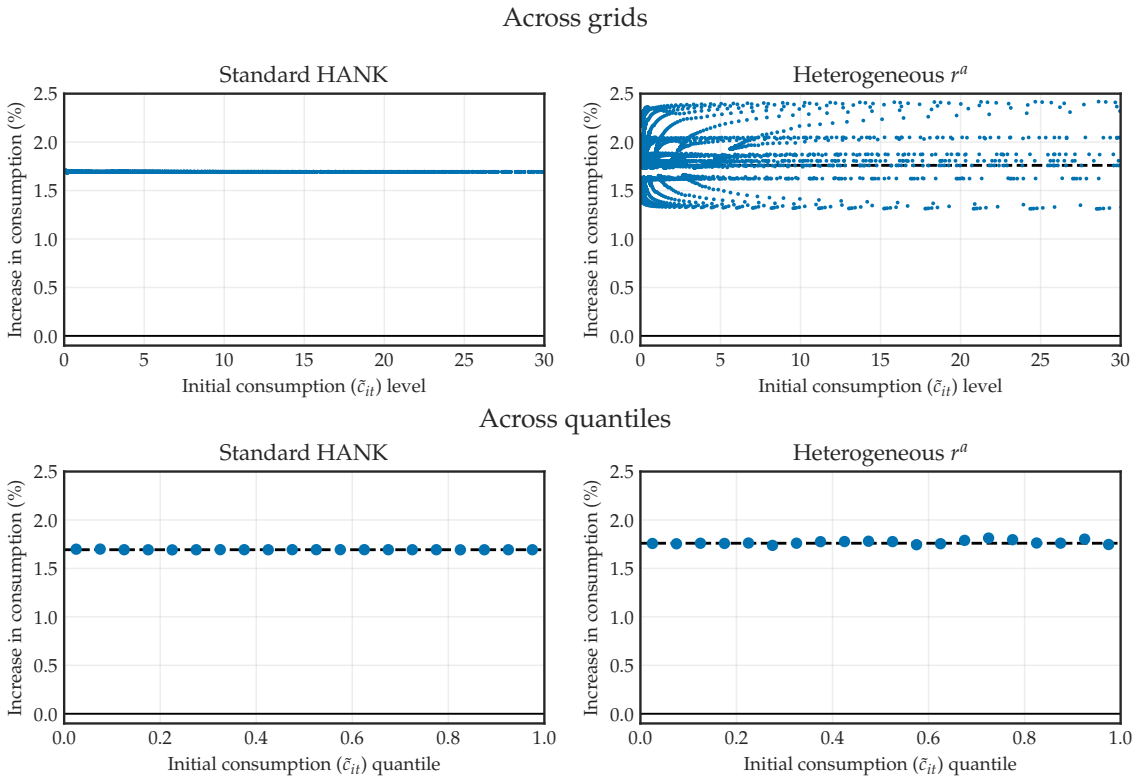


Figure A.23: Heterogeneity in Consumption Response

Note: The figure shows the percent change in consumption to a monetary policy shock, $100(\frac{c_{it}}{\tilde{c}_{it}} - 1)$, at different points in the initial distribution of consumption in two models. The first row shows the response across all grid points. The second row shows the average response within 20 groups based on the initial consumption level.

Table A.9 shows the decomposition of consumption in response to monetary policy, like in Table A.7 in the models consistent with Werning (2015). The table shows that the transmission of monetary policy is very similar in this version of the model, consistent with Figures A.22 and A.23. Additionally, the channels are very similar, which is in contrast with Table A.7. In particular, the redistribution is zero in both models. This is because $\beta^r_{it} = 1$ in this version.

	Standard HANK	Heterogeneous r^a
1. Direct	0.65	0.68
2. Labor	0.23	0.33
3. Taxes	0.00	0.00
4. Revaluation	0.79	0.73
5. Redistribution	0.00	0.00
Total	1.67	1.73

Table A.9: Decomposition of Consumption in Response to Monetary Policy: Werning (2015) Case

Note: The table shows a decomposition of the response of consumption on impact to the monetary policy shock in both models adjusted to be consistent with Werning (2015).

E.4 Equivalent Shocks to Monetary Policy and Fiscal Policy

Let $d\boldsymbol{Y}^{\text{MP}}$ denote the response of output to the monetary policy shock. The response to a government consumption shock is given by

$$d\boldsymbol{Y}^{\text{FP}} = \frac{\partial \boldsymbol{Y}}{\partial \boldsymbol{G}} d\boldsymbol{G},$$

where $\frac{\partial \boldsymbol{Y}}{\partial \boldsymbol{G}}$ is the relevant sequence space general equilibrium Jacobian of output with respect to government consumption. Setting $d\boldsymbol{Y}^{\text{MP}} = d\boldsymbol{Y}^{\text{FP}}$ then gives

$$d\boldsymbol{G} = \left(\frac{\partial \boldsymbol{Y}}{\partial \boldsymbol{G}}\right)^{-1} d\boldsymbol{Y}^{\text{MP}},$$

assuming that the inverse Jacobian exists. The same approach holds for a shock to transfers. The shocks are shown in Figure A.24.

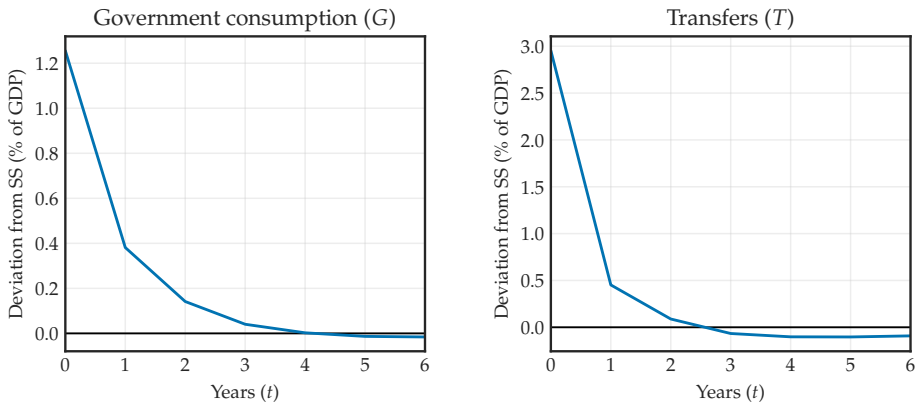


Figure A.24: Fiscal Policy Shocks

Note: The figure shows fiscal policy shocks that induce the same response of output as the monetary policy shock in the model with heterogeneous returns.

E.5 More Micro Effects

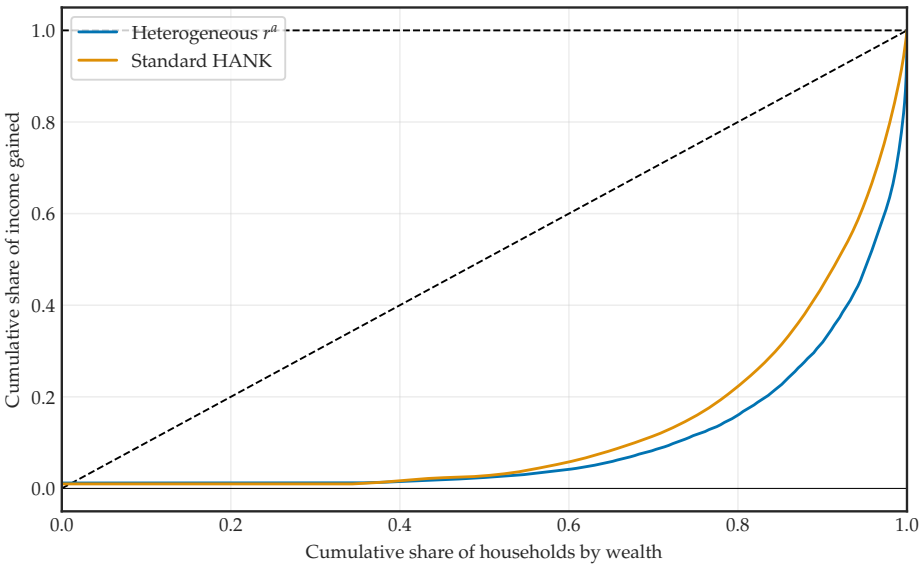


Figure A.25: Distributions of Income Gains In Response To Monetary Policy In Different Models

Note: The figure shows the composition of income generated by monetary policy on impact in two models: The standard HANK model and the HANK model with heterogeneous returns. The plot is smoothed using a Savitzky-Golay filter to smooth out kinks in the solution due to the discretization.

	Heterogeneous r^a	Standard HANK	$\beta^r_{it} = 1$
Transfers	60.3%	61.9%	60.4%
Gov. cons.	21.8%	27.6%	21.9%
Monetary policy	11.3%	17.1%	12.0%

Table A.10: Shares of Income Going to the Bottom 80% In Response

Note: The table shows the shares of income going to the bottom 80% on impact in response to the three different policies in the model with heterogeneous returns but common returns pass-through, $\beta^r_{it} = 1$. "Gov. cons." is government consumption.

	Transfers	Government consumption	Monetary policy
Consumption	−3.5%	0.1%	0.6%
Income	−2.1%	0.4%	1.5%

Table A.11: Effects of shocks on Gini

Note: The table shows the effects of different shocks on the Gini coefficients of different variables. In particular, the figure reports the cumulative (over 6 years) percent change relative to the steady state in the Gini coefficient.

E.6 Income Shares Across Wealth Distribution

Figure A.26 decomposes the income sources from monetary policy along the wealth distribution. Consider first monetary policy. For monetary policy, the bottom 80% of the wealth distribution mainly benefit due to labor income and transfer income, which increase due to general equilibrium effects. At the top of the wealth distribution, the income gain is dominated by capital income, exactly as explained. For transfers and government consumption, the income sources are much more stable along the income distribution, with everyone mainly gaining from the main sources generated by the shock, as shown in Table 4.

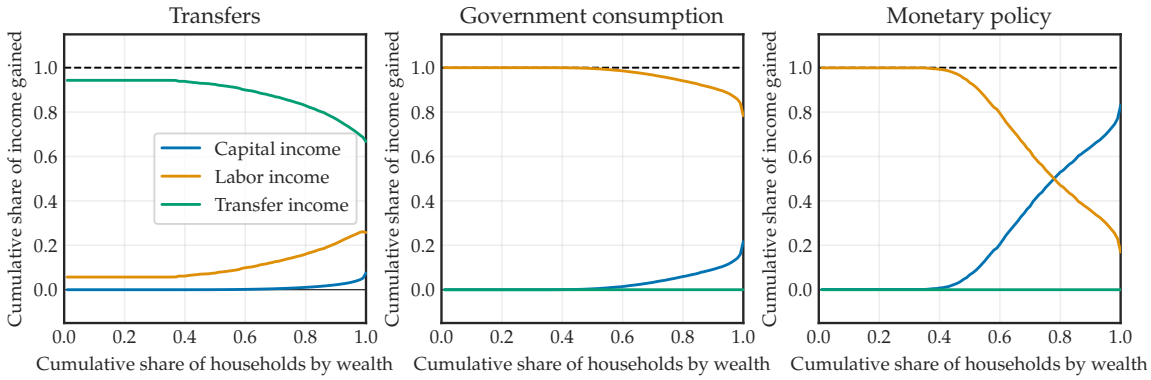


Figure A.26: Income Composition of Policies for Different Households

Note: The figure shows the composition of income generated by different policies on impact along the wealth distribution (a_{it}).

Figure A.27 shows the capital income share due to monetary policy across households in three different models.

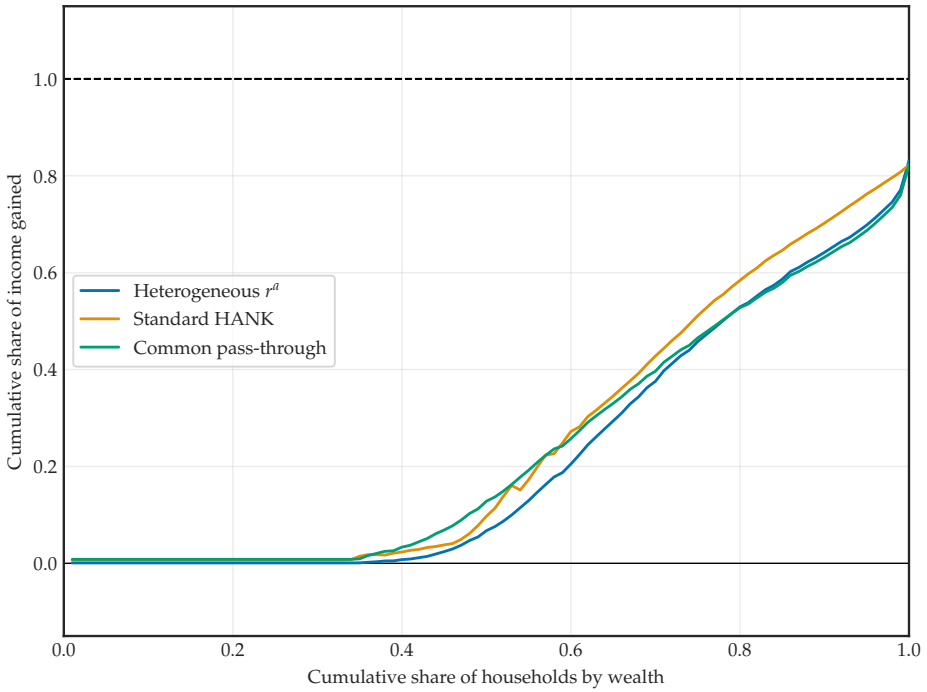


Figure A.27: Capital Income Shares of Monetary Policy In Different Models

Note: The figure shows the share of total income generated by monetary policy on impact due to capital income along the wealth distribution (a_{it}) for three different models: (i) The model with heterogeneous returns, (ii) the standard HANK model, and (iii) the model with heterogeneous returns but common pass-through, $\beta^r_{it} = 1$.

E.7 Comparison to the Empirical Literature

In this Appendix, I compare the distributional effects of monetary policy in my model to the empirical literature. First, I compare to McKay and Wolf (2023a). They estimate the effects of expansionary monetary policy on consumption across the wealth distribution. They find no clear gradient of the consumption response across the wealth distribution. Any gain at the top is largely due to stocks. I do a similar exercise in Figure A.28. The shape of my figure is largely consistent with McKay and Wolf (2023a) and also the model in McKay and Wolf (2023b). I have two takeaways from this. First, my model is consistent with the literature on the consumption response across the wealth distribution. Second, the flat consumption response on impact across coarse buckets of the wealth distribution hides significant heterogeneity in the gains of monetary policy in terms of both income and equivalent variations, particularly at the very top.

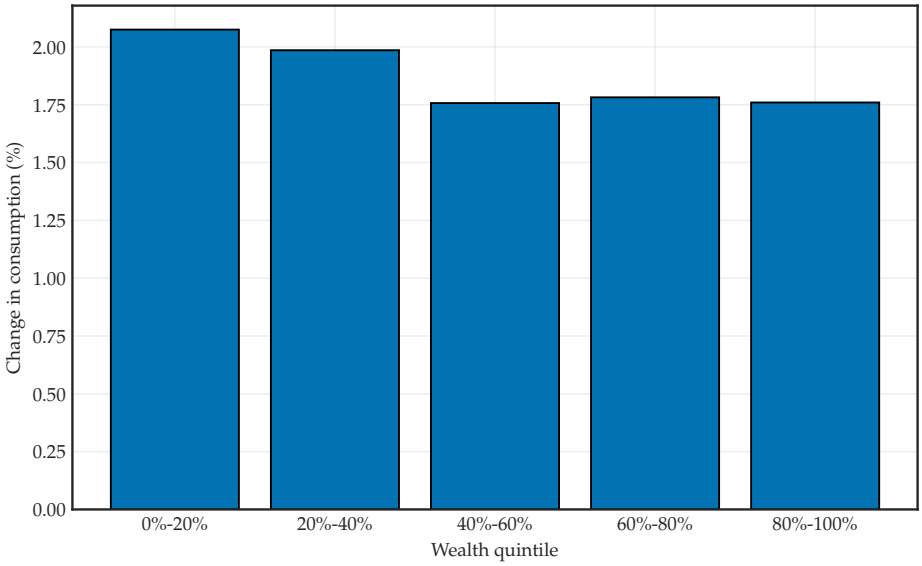


Figure A.28: Consumption Effect Across the Wealth Distribution

Note: The figure shows the percent change in consumption on impact across quintiles of the wealth distribution.

Next, I compare to Andersen et al. (2023), who study the effect of expansionary monetary policy on households in Denmark. Their headline result is a clear income gradient of monetary policy: Households with low income (poor households) lose income, and households with high income (rich households) gain income *relative to the median income change*. The two-year effects range from around a -2% change in income for the poorest and around 4% for the richest, with the median effect being 0%

by definition. I do a similar exercise in Figure A.29. The figure shows a clear income gradient as in Andersen et al. (2023), so my model is broadly consistent with the data.

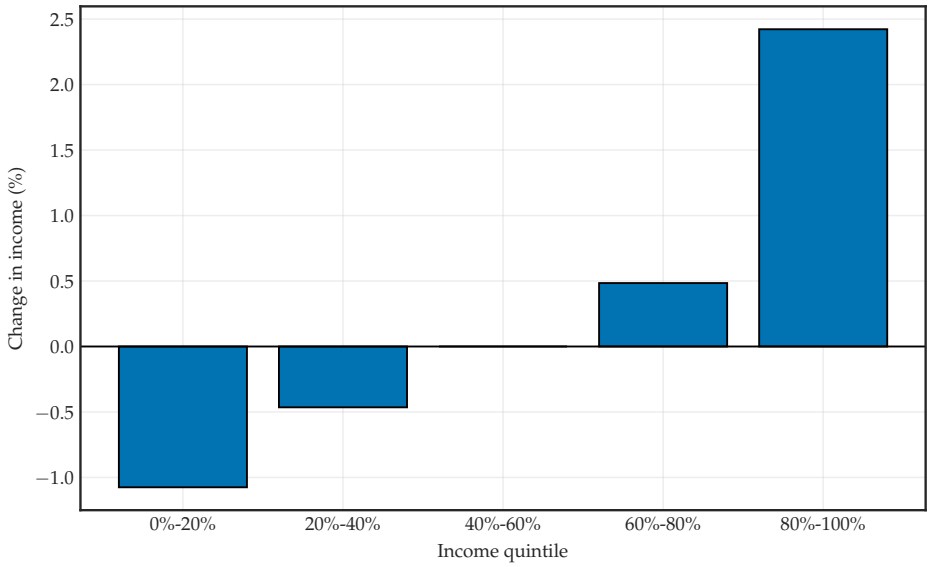


Figure A.29: Income Effect Across the Wealth Distribution

Note: The figure shows the percent change in income on impact across quintiles of the income distribution. The figure is conditioned on income $> 0.1Y_{ss}$ to avoid dividing by income close to 0 when computing percent changes.

E.8 Robustness: Details

E.8.1 Physical Capital

Let me walk through the changes to the model when introducing capital. The production function is now

$$Y_t = \Theta N_t^{1-\alpha} K_{t-1}^{\alpha},$$

where $\Theta > 0$ is a normalization constant chosen such that $Y_{ss} = N_{ss} = 1$, while $\alpha \in [0, 1]$ is the capital share in production. The first order conditions for the firm then read

$$w_t = \frac{1}{\mu}(1-\alpha)\frac{Y_t}{N_t} \quad \text{and} \quad r_t^K = \frac{1}{\mu}\alpha\frac{Y_t}{K_{t-1}},$$

where r_t^K is the real rental rate of capital. Capital is rented from capital firms. Capital firms maximize the discounted sum of profits, facing a virtual adjustment cost, subject

to the law of motion for capital,

$$K_t = (1 - \delta_K)K_{t-1} + I_t,$$

where I_t is investment and $\delta_K \in [0, 1]$ is the depreciation rate. As shown in Druedahl et al. (2025), this yields the following first-order conditions:

$$1 + \phi_I \left(\frac{I_t}{I_{t-1}} - 1\right) \frac{I_t}{I_{t-1}} + \frac{\phi_I}{2} \left(\frac{I_t}{I_{t-1}} - 1\right)^2 = Q_t + \frac{1}{1 + r_t} \phi_I \left(\frac{I_{t+1}}{I_t} - 1\right) \left(\frac{I_{t+1}}{I_t}\right)^2,$$
$$Q_t = \frac{1}{1 + r_t} \left[(1 - \delta_K)Q_{t+1} + r^K_{t+1}\right],$$

where $\phi_I \geq 0$ controls the adjustment cost and Q_t is Tobin's Q. Profits from the capital firms are rebated to the final goods producers, so

$$D_t = Y_t - w_t N_t - I_t.$$

Goods market clearing then finally reads

$$Y_t = C_t + G_t + I_t.$$

I set $\alpha = 1/3$, $\delta_K = 0.05$, and $\phi_I = 9.6$ following Auclert et al. (2020).

E.8.2 Long-Term Debt

I consider what happens if government bonds have a duration longer than 1 year. In particular, the government issues a quantity of real bonds, B_t, with real price, q_t. The budget constraint in eq. (8) is then modified to

$$q_t B_t = (1 + \delta q_t)B_{t-1} + G_t + T_t - \mathcal{T}_t. \tag{19}$$

The bonds are long, paying a unit coupon each period. They are exponentially decaying with decay rate $\delta \in [0, 1]$, cf. Auclert and Rognlie (2018) and Auclert et al. (2020). I consider the case of long bonds because the literature emphasizes the importance of this for the transmission of monetary policy, cf. Auclert (2019).

In addition to these changes, the definition of capital income in eq. (13) becomes

$$\underbrace{\int r^a_{it} a_{it-1}\, di}_{\text{HH's capital income}} = \underbrace{p_t + D_t + (1 + \delta q_t)B_{t-1} - A_{t-1}}_{\text{Capital income from bonds and equity}}. \tag{20}$$

Finally, asset market clearing reads

$$A_t = p_t + q_t B_t.$$

For the calibration, I set $\delta = 0.8$ as in Auclert et al. (2020) to match the average US debt maturity of 5 years.

E.8.3 Nominal Debt

I now consider what happens if government bonds are nominal. In this case, the government's budget constraint in eq. (8) is modified to

$$B_t = \frac{1 + i_{t-1}}{1 + \pi_t} B_{t-1} + G_t + T_t - \mathcal{T}_t. \tag{21}$$

The definition of capital income in eq. (13) reads

$$\underbrace{\int r_{it}^a a_{it-1}\, di}_{\text{HH's capital income}} = \underbrace{p_t + D_t + \frac{1 + i_{t-1}}{1 + \pi_t} B_{t-1} - A_{t-1}}_{\text{Capital income from bonds and equity}}. \tag{22}$$

E.8.4 Sticky Prices

I now consider what happens if the firm cannot set the price every period but instead has sticky prices. In particular, firms set the price, P_t, subject to quadratic adjustment costs, with discount factor $(1 + r_t)^{-1}$. This yields the following new Keynesian Phillips curve (NKPC) for inflation, $\pi_t = P_t / P_{t-1} - 1$, which replaces eq. (7):

$$\log(1 + \pi_t) = \kappa^P \left(w_t - \frac{1}{\mu} \right) + \frac{1}{1 + r_t} \frac{Y_{t+1}}{Y_t} \log(1 + \pi_{t+1}).$$

Here, $\kappa^P \geq 0$ is the slope of the Phillips curve. Flexible prices are obtained as $\kappa^P \to \infty$. I set $\kappa^P = 0.23$ as in Auclert et al. (2024b). Crucially, this implies $\kappa^W < \kappa^P$ such that wages are more sticky than prices (Broer et al. 2020).

E.8.5 Parametrization

In addition to the model changes, I also consider three different parametrizations of the model. First, I consider what happens when wages are more flexible: $\kappa^W = 0.1$ instead of $\kappa^W = 0.03$. Second, I consider the case of more liquidity being in the form of government bonds: $B/A = 50\%$ instead of the baseline of $B/A = 23\%$. This implies a different markup, μ, which is important. To see this, note that the gains of the rich in the model with heterogeneous returns largely go through capital income.

But what drives capital income in the model? And is the magnitude of the capital income gains empirically realistic? The key to answering this question is Proposition 3, which decomposes the effects of aggregate shocks on capital income.

Proposition 3. *The effect of an aggregate shock on capital income can be decomposed into two components: The firm equity price and dividends. This decomposition can be written as follows:*

$$dx_0 = dp_0 + dD_0. \tag{23}$$

These components are given by:

$$D_0 = \frac{\mu - 1}{\mu} Y_0,$$
$$p_0 = \frac{1}{1 + r_0} \frac{\mu - 1}{\mu} Y_1 + \frac{1}{(1 + r_0)(1 + r_1)} \frac{\mu - 1}{\mu} Y_2 + \dots$$

Third, I consider adding net borrowing to the model. This amounts to changing the borrowing constraint from $a_{it} \geq 0$ to $a_{it} \geq -0.25 Y_{ss}$. I chose $0.25 Y_{ss}$ motivated by Kaplan et al. (2018).

F Appendix to Section 7

F.1 Decomposing Equivalent Variations

In this Appendix, I decompose the equivalent variations in Figure 16. I do this by noting that the equivalent variation of household i is (only) a function of the sequences of aggregate inputs to the household problem and labor supply: r_t^a, Z_t, T_t, τ_t, and N_t. In general, a shock will change these sequences compared to the steady state, resulting in the equivalent variations. To decompose this equivalent variation, I set all sequences to their steady state values except one at a time. This gives the contribution of that variable to the equivalent variation. Summing across all variables should approximately equal the actual equivalent variation, but this does not hold precisely due to non-linearities.

Figure A.30 shows the decompositions of the three shocks. The channels are as follows: Revaluation (changing r_0^a), earnings (changing Z_t), transfers (changing T_t), and everything else (changing N_t, τ_t, and r_t^a for $t > 0$). The channels are as a share of the total, i.e., they sum to 1.



Figure A.30: Decomposition of Equivalent Variations

Note: The figure shows a decomposition of equivalent variations into channels in response to different shocks. The plot is smoothed using a Savitzky-Golay filter to smooth out kinks in the solution due to the discretization.

F.2 Estimation Details

In this Appendix, I present the estimation of the model. The point of this exercise is simply to have some exogenous variation in business cycles that approximates business cycles in the data for the policymaker to stabilize. To do so, I introduce three shocks to the model:

1. **Markup shock.** The markup, μ, is replaced by $\mu_t = \mu + \epsilon_t^{\mu}$, where ϵ_t^{μ} is a shock.
2. **Risk premium shock.** The no-arbitrage condition is replaced by

$$p_t = \frac{p_{t+1} + D_{t+1}}{1 + r_t + \epsilon_t^{\text{rp}}},$$

   where ϵ_t^{rp} is a shock.
3. **NKWPC shock.** The NKWPC is replaced by

$$\pi_t^W(1 + \pi_t^W) = \kappa^W \left(\frac{v'(N_t)}{u'(C_t)(1 - \tau_t)w_t} - 1 \right) + \frac{1}{1 + r_t} \pi_{t+1}^W (1 + \pi_{t+1}^W) + \epsilon_t^{\text{pc}},$$

   where ϵ_t^{pc} is a shock.

Let me briefly discuss why I chose these shocks. Essentially, these are all the standard shocks in the model that make sense for my purpose. In particular, I rule out the following types of shocks for my purpose that one could otherwise consider. I rule out TFP shocks as they do not make much sense with linear production. I rule out any policy shocks because the point is to set policy. I rule out shocks to preferences (like discount factor or risk aversion shocks) because the point is to do a welfare analysis.

For estimation, I assume that each of the three shocks follow AR(1) processes with IID normal innovations:

$$\epsilon_t^x = \rho_x \epsilon_{t-1}^x + \varepsilon_t^x,$$

where $\varepsilon_t^x \sim \text{iid} N(0, \sigma_x^2)$. I estimate these parameters, i.e., I estimate the vector

$$\theta = (\rho_\mu, \rho_{\text{rp}}, \rho_{\text{pc}}, \sigma_\mu, \sigma_{\text{rp}}, \sigma_{\text{pc}}).$$

To do so, I use maximum likelihood estimation (MLE). In practice, I use the method and software implementation provided in Auclert et al. (2021). I do this by assuming a passive policymaker, i.e., by setting $r_t = r_{ss}$, $G_t = G_{ss}$, and $T_t = T_{ss}$. The question is then: What parameters, θ, make the model produce a simulation most similar to the data?

For the data, I use two time series for the US for 1950–2019: Consumption and inflation. The data is from FRED and is as follows.

- **Consumption.** I start with FRED series "PCECCA", which is annual real consumption in chained dollars. To obtain dC_t, I first estimate a regression of log

consumption on $(1, t, t^2)$, where t is the year. I use the exponential of the predicted log consumption as the consumption trend. I then compute dC_t as the relative deviation of consumption from its trend.

- **Inflation.** I start with FRED series "A191RD3A086NBEA", which is the annual GDP implicit price deflator. I then calculate inflation as the percent change in the price level. To obtain $d\pi_t$, I demean this inflation series.

The time series are plotted in Figure A.31.

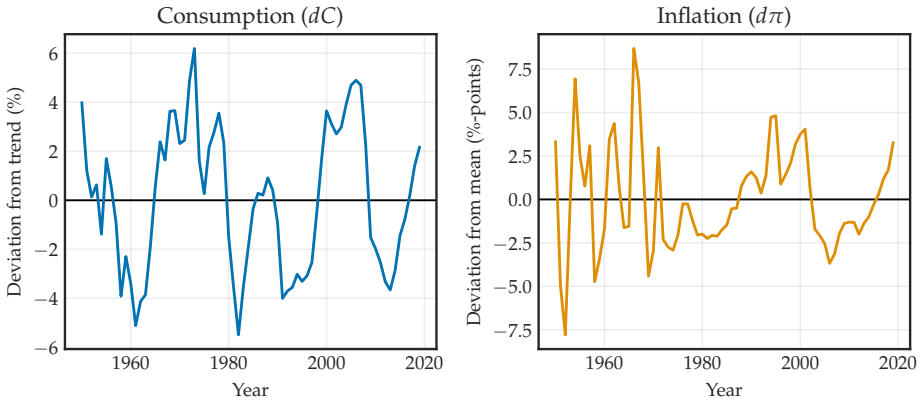


Figure A.31: Data Time Series

Note: The figure shows de-trended consumption, dC_t, and inflation, $d\pi_t$, in the US. These are the time series used for estimating the model.

The estimates are reported in Table A.12.

NKWPC, ϵ^{pc}		Markup, ϵ^{μ}		Risk premium, ϵ^{rp}	
ρ_{pc}	σ_{pc}	ρ_{μ}	σ_{μ}	ρ_{rp}	σ_{rp}
0.73	0.63	0.87	0.86	0.00	0.34

Table A.12: Estimation Results

Note: The table shows estimation results.

F.3 Simulation Methodology

Consider simulating some variable, X_t, for a time series $t = 0, 1, \ldots, T-1$, in response to aggregate shocks. This could be consumption, consumption of the rich, labor, etc.

To first order, the time series is given by

$$dX_t = \sum_j \sum_{s=0}^{\infty} \text{IRF}_{j,s} \varepsilon^j_{t-s}.$$

Here, j indexes the shock, ε^j_t, and $\text{IRF}_{j,s}$ is the IRF of X_t to a shock that occurred s periods ago. I discard a burn-in period of length 1000. I first do this in a model where monetary policy is used to stabilize consumption. I then do it in a model where fiscal policy is used to stabilize consumption. Denote the IRFs in the two cases by $\text{IRF}^{\text{MP}}_{j,s}$ and $\text{IRF}^{\text{FP}}_{j,s}$. Consider then the asymmetric policy where fiscal policy is used to stabilize shocks that create recessions and monetary policy to stabilize shocks that create booms. Thus,

$$dX^{\text{As.}}_t = \sum_j \sum_{s=0}^{\infty} \text{IRF}^{\text{FP}}_{j,s} \mathbb{1}(\varepsilon^j_{t-s} > 0) \varepsilon^j_{t-s} + \sum_j \sum_{s=0}^{\infty} \text{IRF}^{\text{MP}}_{j,s} \mathbb{1}(\varepsilon^j_{t-s} < 0) \varepsilon^j_{t-s}.$$

The idea here is that positive shocks, $\varepsilon^j_t > 0$, create a recession. Technically, the policymaker therefore does not respond to the recession itself, but to the shock that creates the recession.[34]

F.4 Simulation Robustness

Version	Heterogeneous r^a	Standard HANK
Cutoff = P80	1.62 pp.	1.55 pp.
Cutoff = P90	1.46 pp.	1.33 pp.
Cutoff = P95	1.29 pp.	1.10 pp.
Cutoff = P99	1.06 pp.	0.75 pp.

Table A.13: Difference of Welfare Gains For Poor and Rich

Note: The table shows the difference in the welfare gains between the poor and the rich in the case of asymmetric policy using fiscal policy in recessions and monetary policy in booms. The rows refer to the percentile of wealth determining the cutoff between rich and poor. The columns refer to the models.

34. The exception is the NKWPC shock, which does not affect output, so policy does nothing.

F.5 Simulation Results

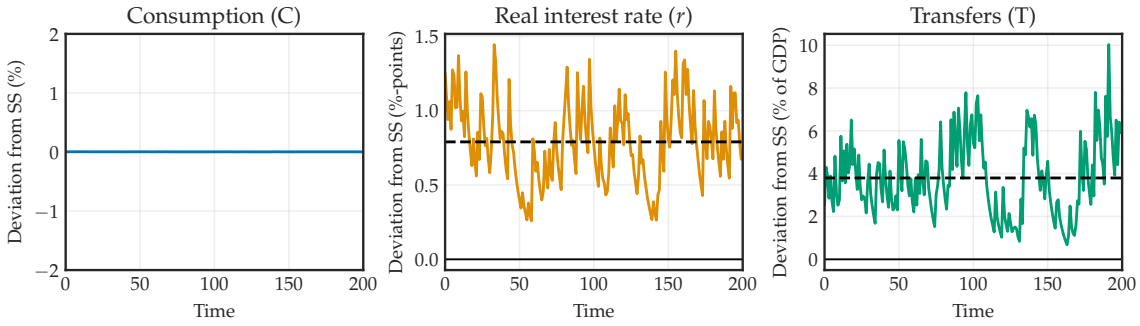


Figure A.32: Simulated Time Series

Note: The figure shows 200 periods of the simulated time series with an asymmetric policy using transfers to stabilize shocks that create recessions and monetary policy to stabilize shocks that create booms.