

HANK Comes of Age

Bence Bardóczy ^{*1} and Mateo Velásquez-Giraldo ^{†2}

¹Federal Reserve Board

²Johns Hopkins University

June 21, 2024

Abstract

We study the aggregate and distributional effects of monetary policy in a heterogeneous agent New Keynesian model that explicitly represents the life cycle of households. The model matches the age patterns in the level and dispersion of labor income and financial wealth in the U.S. despite the absence of preference heterogeneity and portfolio adjustment costs. Monetary policy affects the consumption of young households mainly through labor income and the consumption of old households mainly through asset returns. More than half of the aggregate consumption response to an expansionary monetary policy shock comes from those below the age of 40. The shock redistributes welfare from the wealthiest old to the poorest young and increases average welfare of most cohorts.

*bence.a.bardoczy@frb.gov

†mvelasq2@jh.edu

The views expressed in this paper are those of the authors and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff.

We thank Christopher Carroll, Hess Chung, and Eva F. Janssens for helpful comments. Harrison Snell provided excellent research assistance.

1 Introduction

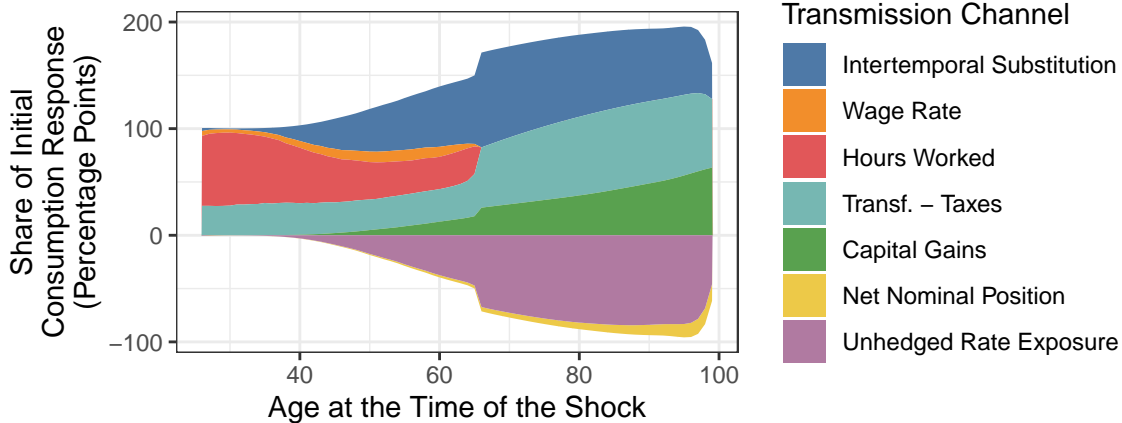
A growing literature has convincingly argued that the transmission of monetary policy to macroeconomic aggregates works in large part through mechanisms that representative-agent models omit or subdue. Households, for instance, have different levels of savings allocated to different assets, different work arrangements with different sensitivities to the business cycle, and different propensities to consume out of changes in their income and their wealth. When these dimensions of heterogeneity are exposed to general-equilibrium changes in labor markets, fiscal policy, and asset prices, they produce indirect effects that, in the case of consumption, can be greater than those of traditional mechanisms like intertemporal substitution. While households experience large changes in all of these dimensions of heterogeneity as they age, the heterogeneous agent New Keynesian (HANK) models used to study indirect transmission channels feature infinite-horizon or perpetual youth frameworks that omit these changes. An explicit treatment of the life cycle remains absent from this class of models.

A more realistic treatment of the life cycle of households in the HANK models that inform macroeconomic policymaking could bring several benefits. First, this enhancement would allow more detailed and realistic assessments of the heterogeneous effects of macroeconomic policies on consumption, wealth, and welfare across households of different ages. The relevance of these assessments is supported by abundant empirical evidence showing that households of different ages differ in their responses and exposures to the transmission channels of monetary policy.¹ Second, since age is measured in most micro level data sources, these assessments would constitute a greater set of testable predictions against which these models could be evaluated. Testing these predictions can advance our understanding of both the effects of macroeconomic policies on economic activity and the factors and conditions that modulate them. Third, by sidestepping life-cycle dynamics, current models omit important economic considerations like the need to save for retirement. Incorporating these considerations can improve upon the known limitations of these models.²

In this paper, we embed a realistic life-cycle model of households into a state-of-the-art New Keynesian framework to study how the transmission channels and welfare effects of monetary policy vary across households of different ages. The representation of households is realistic in the sense that it matches the life-cycle patterns in the level and dispersion of their income and wealth. The New Keynesian framework captures the dynamic general-equilibrium effects that monetary policy produces on asset and goods prices, labor markets, and fiscal policy. The model of households captures the individual changes in consumption, wealth, and welfare stemming from aggregate fluctuations, taking into account idiosyncratic

¹For example, Doepke and Schneider (2006), Adam and Zhu (2016), Greenwald et al. (2022), Pallotti et al. (2023), and Fagereng et al. (2023) find heterogeneous exposure to inflation, asset prices, and interest rates across the life cycle. A large literature also documents that the exposure to labor-market fluctuations varies across the life cycle (Clark & Summers, 1981; Sabelhaus & Song, 2010; Jaimovich et al., 2013; Guvenen et al., 2017).

²For instance, the omission of the life-cycle and bequest motives for saving could be a reason behind the inability of current single-asset models to simultaneously match broad measures of wealth holdings and sizable marginal propensities to consume. Indeed, our results favor this hypothesis.



This figure shows a decomposition of the monetary transmission mechanism by age. We consider an expansionary shock as described in Section 4. We index cohorts by their age at the time of the shock. We calculate the on-impact change in the total consumption of each cohort. To decompose the response into its channels, we calculate the responses to the shock-induced changes in subsets of variables (for example, wages, on-impact asset returns, taxes, and transfers), leaving all others in their steady state-values. Most channels are self-explanatory; see Section 4 for details. “Net Nominal Position” represents a scenario where only the initial return in bond returns, R_t^b , is passed to the household block. “Capital Gains” isolates the effect of initial equity revaluation, only the initial return to stocks R_t^s changes. “Unhedged Rate Exposure” inputs the realized return changes after initial revaluations $\{R_{t+s}^b, R_{t+s}^s\}_{s \geq 1}$.

Figure 1: Transmission of an Expansionary Monetary Policy Shock Across the Life Cycle

uncertainty and heterogeneity by age, income, and wealth. Combining the two allows us to quantify the total effect of monetary policy on households of different ages and the relative contributions of different channels, both direct and indirect. As a preview of our results, Figure 1 decomposes the consumption response to a monetary easing into its transmission channels, showing that their relative contributions vary over the life cycle.

The main features of our household block are an age-varying stochastic income process, a bequest motive that becomes more relevant for wealthier households, and borrowing constraints. Households save to protect their consumption against income shocks, increase their consumption during retirement, and leave bequests. The three motives turn out to be important for consumption dynamics. Uninsurable shocks and borrowing constraints generate marginal propensities to consume (MPCs) of magnitudes comparable with those of empirical estimates. Life-cycle dynamics help us match a broad measure of wealth—total financial assets—while preserving a large annual aggregate MPC of 0.41. This has been a challenge for one-asset HANK models (see Kaplan & Violante, 2022), and we achieve it by allowing young households with high MPCs to coexist with wealthy retirees. Finally, bequests make most retired households run down their assets slowly (or not at all), which is consistent with empirical findings on the savings of the old, and lowers their MPCs.

The other elements of our general equilibrium model have structures that are standard in the HANK literature. Production uses capital and labor and features quadratic costs of changing the price of goods and the stock of capital. Shares of the firms are tradable and owned by households and the government, who receive their dividends. Nominal wages are set by a labor union, also with adjustment costs. The monetary authority follows an

inflation-targeting Taylor rule, and the fiscal authority adjusts a progressive tax function to maintain a balanced budget. These elements generate equilibrium dynamics of inflation, investment, asset prices, taxation, and labor markets, all of which affect the household block.

We study the effects of an expansionary monetary policy shock and the mechanisms behind them through the lens of our model. Like previous studies, we find that indirect effects are responsible for the lion’s share of the response in aggregate consumption; we contribute new results on the incidence of the total response and its mechanisms across households of different ages, which turn out to vary substantially. For instance, households aged 40 and below constitute less than 30% of the population but account for almost 60% of the initial response of aggregate consumption. As Figure 1 shows, their response comes almost exclusively from changes in their after-tax labor income. For retirees, changes in present and future asset returns play a bigger role: intertemporal substitution becomes the main driver of the response, and there are substantial wealth-revaluation effects.

We study the welfare effects of the expansionary shock and find that, while most cohorts benefit from it, the average welfare benefit is much greater for those who are younger at the time of the shock.³ However, there are wide differences within age groups, as the welfare impacts also have a steep relationship with wealth. These two dimensions of heterogeneity combine to reveal a sizable welfare redistribution from the old and wealthy to the young and poor: the losses of retirees in the highest wealth quintile are comparable to 3.6% of their one-year consumption, while the gains of prime-age households in the lowest wealth quintile are greater than 7.0% of theirs. This pattern is the combination of labor-income effects that are positive for most households (especially the poorest) and asset-returns effects that are small for most households except those above the age of 45 and in the two highest wealth quintiles. For the latter, the negative effect of asset returns can be as large as 7.0% of their consumption in spite of a revaluation of their real assets at the time of the shock. The reason is that the future sequence of lower real returns that accompanies the expansionary shock has a negative effect on the feasibility of these households’ desired consumption-savings paths—they have unhedged exposure to the real interest rate (Auclert, 2019).

Related literature. Our paper relates to four main groups of studies in the macroeconomics and household-finance literatures.

The first group of related papers has focused on reproducing particular features of the distribution of wealth in the United States, paying special attention to the fact that wealth is distributed more unevenly than income. Standard models with homogeneous preferences in which households accumulate wealth primarily for precautionary reasons struggle to generate the level of inequality observed in U.S. data (Quadrini & Rios-Rull, 1997; De Nardi, 2015). Studies like Benhabib and Bisin (2018) and Stachurski and Toda (2019) theoretically dissect model properties that may and may not yield the observed level of wealth inequality and its relationship to income inequality. Among the model properties that have been tried,

³We maintain the assumption that long-run inflation expectations are anchored to the central bank’s target. Taking the risk of de-anchoring into account may lower the broadly positive welfare effects of monetary expansion.

there are earnings processes with “awesome” or “superstar” (very-high-income) states that are calibrated directly to match features of the wealth distribution (for example, Castañeda et al., 2003); heterogeneous time-discount factors (for example, Carroll et al., 2017); non-homothetic preferences that generate saving rates that increase with wealth (for example, Carroll, 2002; De Nardi, 2004); and heterogeneous returns to wealth (for example, Benhabib et al., 2019). Our approach to generating wealth inequality combines non-homothetic preferences in the form of a “luxurious” bequest motive and a skewed income process. Among the available approaches, the bequest motive allows us to match the fact that some old households run down their wealth slowly or not at all. Although our model does feature the chance of entering a state with very high income, we do not estimate any feature of the *income* process to match the distribution of *wealth*; our income process comes from Arellano et al. (2017), who estimate it using income data alone. For a model with homogeneous preferences and returns, and income estimates that do not target wealth data, our fit of the Lorenz curve for wealth is remarkable.

The second group of papers are those in the growing literature that uses HANK models to study monetary policy and its transmission. Early contributions to this literature include McKay et al. (2016), Guerrieri and Lorenzoni (2017), Kaplan et al. (2018), Bilbiie (2018), and Auclert (2019). Our findings regarding the aggregate effects of monetary policy and its mechanisms qualitatively align with theirs—indirect transmission channels are of chief importance, because a significant share of the population has a high MPC. An important difference with these studies is that our model generates these high MPCs while matching households’ total financial wealth; past modeling efforts have only been able to generate high MPCs if they target narrow measures of wealth or model a significant fraction of this wealth as illiquid (Kaplan & Violante, 2022). Our inclusion of life-cycle and bequest motives for saving are behind this achievement.

The third group of papers are those that study monetary policy in economies with overlapping generations. Braun and Ikeda (2021), Bielecki et al. (2022), and Bullard et al. (2023) are in this small group. While these studies incorporate representations of life-cycle variation in the income and assets of households, they abstract away from within-cohort heterogeneity or limit it to a small number of ex-ante types. These simplifications carry implausible implications for MPCs—which are crucial to the study of indirect channels (Auclert, 2019)—and for welfare.⁴ Our paper is, to the best of our knowledge, the first to study monetary policy and its transmission channels in a model that features both overlapping generations and within-cohort heterogeneity that is not limited to ex-ante types but arises also from uninsurable shocks. We find that both between-cohort and within-cohort heterogeneity are substantial: there are large differences in the consumption responses, welfare impacts, and transmission mechanisms of monetary policy across both dimensions.

The fourth and final group of papers that we relate to studies the distributional consequences of changes in inflation, real interest rates, and asset prices. Doepke and Schneider

⁴Both the consumption and welfare functions are nonlinear. The change (induced by, say, a transfer) in the consumption and welfare of a household that owns the average wealth can be substantially different to the average change in the consumption and welfare computed over households that own their actual wealth.

(2006), Adam and Zhu (2016), Greenwald et al. (2022), Pallotti et al. (2023), and Fagereng et al. (2023) are in this category. All of these papers find that age is a prominent dimension along which there is redistribution when inflation, interest rates, or asset prices change: households of different ages hold different amounts of assets and liabilities, of different types (real or nominal), and of different duration. Important subtleties arise when evaluating these redistributions. For example, with lower interest rates, a household whose wealth initially increases due to the repricing of real assets may still find that it can no longer afford its original consumption plan. Hence, wealth and welfare may not move in lockstep (Auclert, 2019; Greenwald et al., 2022; Fagereng et al., 2023). Our model accounts for these subtleties, delivering a measure of welfare that encompasses initial revaluations, dynamic considerations, and the optimal reaction of households to their new conditions. We study the redistribution of welfare generated by expansionary monetary policy shocks and the channels through which it operates.

The rest of this paper is organized as follows. Section 2 presents our life-cycle model of households, its calibration, and its implications about the distribution of income, wealth, and MPCs. Section 3 discusses all the other blocks of our New Keynesian model. In Section 4, we study the response of the model economy to an expansionary monetary policy shock, examining transmission mechanisms, heterogeneous responses, and welfare redistribution. Section 5 concludes.

2 Life-Cycle Model of Households

Households are born at age 26 and live up to a maximum age of 100, facing a stochastic probability of death every year.⁵ Every period, they decide how much to consume out of their income and accumulated assets. They may save their resources to insure against income shocks, to prepare for retirement, and to leave bequests. This section describes the various elements of their intertemporal problem and how we calibrate them. Throughout the section, we index individuals with i and time with t .

2.1 Income Process

Agents work and earn market wages until the age of 65. After that, they retire and start receiving income flows that represent Social Security benefits and pensions.

Let $\mathbf{a}_{i,t}$ denote the age of individual i at time t . Let $y_{i,t}$ be pretax income, w_t be the prevailing wage, $\tilde{y}_{i,t}$ be the endowment of efficiency units, and $n_{i,t}$ be hours worked. We adopt the following specification for income in working years ($\mathbf{a}_{i,t} \leq 65$)

$$\begin{aligned} y_{i,t} &= \tilde{y}_{i,t} \times w_t \times n_{i,t} \\ \ln \tilde{y}_{i,t} &= \alpha_i + f_{\mathbf{a}_{i,t}} + z_{i,t} \\ z_{i,t+1} &\sim \Pi_{\mathbf{a}_{i,t}}(z_{i,t}) \end{aligned} \tag{1}$$

⁵We use age-specific death probabilities from the SSA life-tables. Figure C.1 in Appendix C shows the age distribution of our simulated populations.

and the following for income in retirement years ($\mathbf{a}_{i,t} > 65$)

$$\begin{aligned} y_{i,t} &= d_{i,t} \\ \ln d_{i,t} &= \alpha_i + f_{\mathbf{a}_{i,t}} + z_{i,t(\mathbf{a}_{i,t}=65)}. \end{aligned} \tag{2}$$

There are individual fixed-effects α_i and age fixed-effects $f_{\mathbf{a}}$ for productivity. The final component of productivity, $z_{i,t}$, is a persistent shock. This shock follows a Markov process with age-specific support and transition probabilities during working years and then shuts off in retirement. Retirement benefits $d_{i,t}$ are paid by the government. We assume that retirement benefits scale with the value of the persistent shock in the last working year of the agent, $z_{i,t(\mathbf{a}_{i,t}=65)}$.⁶

For the persistent shock $z_{i,t}$, we use the age-specific persistent income shock process of Arellano et al. (2017) provided by Janssens and McCrary (2023)⁷. Given the critical role of income shocks for the determination of savings and MPCs, we discuss the key features of the income process in appendix A.1.

2.2 Hours and Labor Demand

Hours worked are determined by labor demand and not chosen by individual households. Firms demand a total amount of productivity-adjusted hours, L_t . In steady state, we normalize the hours worked by individual households to $n_{i,ss} = 1$, which implies that L_{ss} is equal to the average productivity of working-age households. When aggregate labor demand deviates from steady state, individual hours $n_{i,t}$ must adjust to maintain

$$L_t = \int \tilde{y}_{i,t} \times n_{i,t} dD_t(\mathbf{a}_{i,t} \leq 65). \tag{3}$$

How to pin down $n_{i,t}$ across households? It is a well-established fact that the hours of young workers fluctuate more over the business cycle than that of older workers (see, for example, Clark & Summers, 1981; Gomme et al., 2004; Jaimovich et al., 2013).⁸ Jaimovich et al. (2013) argue that the higher volatility in the hours of the young comes mostly from labor demand, not supply, and rationalize this fact on the basis of capital-experience complementarities in production.

In light of the evidence, we allow the adjustments in hours necessary to maintain (3) to be unevenly borne by households of different ages. Specifically, we assume that individual hours are a function of age and aggregate labor demand $n_{i,t} = \gamma(\mathbf{a}_{i,t}, L_t)$ and adopt a flexible function that guarantees that total effective hours aggregate to total demand. Letting $\tilde{Y}_{\mathbf{a}} \equiv$

⁶This is a common modeling device that allows for retirement benefits that scale with income without having to track the average lifetime income of the agent as an additional state variable. See for example Carroll (1997) and Kaplan and Violante (2014).

⁷We thank Eva Janssens for providing us with the optimally discretized income process.

⁸Sabelhaus and Song (2010) and Guvenen et al. (2017) demonstrate a similar fact for fluctuations in earnings.

$\int_{i:\mathbf{a}_i=\mathbf{a}} \tilde{y}_{i,t}$ be the total effective units per hour available from individuals of age \mathbf{a} , we set

$$\gamma(\mathbf{a}_{i,t}, L_t) = L_t \times \frac{\left(\frac{L_t}{L_{ss}}\right)^{\varepsilon_{\mathbf{a}_{i,t}}}}{\sum_{\mathbf{a}} \tilde{Y}_{\mathbf{a}} \times \left(\frac{L_t}{L_{ss}}\right)^{\varepsilon_{\mathbf{a}}}}, \quad (4)$$

where $\varepsilon_{\mathbf{a}}$ are parameters that control the sensitivity of age- \mathbf{a} demand to aggregate demand.

We calibrate the hours function to match the ratio of age-specific hour volatilities to aggregate hours volatility reported by Jaimovich et al. (2013). Appendix A.2 discusses our targets and calibration strategy. Figure A.3 displays our estimated function for different levels of aggregate demand, showing that the sensitivity of hours to aggregate demand decreases with age.

2.3 Assets

Households save using two different assets: government bonds and firm equity. Bonds and equity have return factors R^b and R^s , respectively. Because there is no aggregate uncertainty, every agent expects (demands) that $R^b = R^s$ both in steady state and along perfect foresight transitions. However, unexpected shocks—“MIT shocks”—of the kind that we analyze in this paper can generate returns for bonds and stocks that differ ex-post since, at the time of their announcement, the price of equity jumps to be consistent with the newly announced conditions.

Because they expect the return factors of bonds and equity claims to be identical and deterministic, individuals are indifferent about the allocation of their assets. Their portfolio shares are not determined by their optimizing behavior. Therefore, we set portfolio shares exogenously to match their relationship with age and wealth in the 2019 Survey of Consumer Finances (SCF), and the aggregate firm ownership by households that is required to clear markets in general equilibrium. The share of assets that household i invests in equity claims is a function of its age and its assets, $\zeta(\mathbf{a}_{i,t}, a_{i,t-1})$. The aggregate share of household assets invested in equity, which we denote with Ξ_t , is

$$\Xi_t = \frac{\int \zeta(\mathbf{a}_{i,t}, a_{i,t-1}) \times a_{i,t-1} dD_t}{\int a_{i,t-1} dD_t}. \quad (5)$$

We calibrate the function $\zeta(\mathbf{a}_{i,t}, a_{i,t-1})$ to match a required steady-state aggregate share Ξ_{ss} and also to be consistent with the equity-holding patterns across age and wealth in the 2019 SCF. Appendix A.3 depicts $\zeta(\cdot, \cdot)$ and describes how we estimate it.

2.4 Taxation

Individuals pay income taxes both during their working years and in retirement. We use the functional form proposed by Heathcote et al. (2017), which we denote with a function \mathcal{T} , that transform pretax into after-tax income. The function is

$$\mathcal{T}(y_{i,t}) = \lambda_t y_{i,t}^{1-\tau}, \quad (6)$$

so that taxes are $y_{i,t} - \lambda_t y_{i,t}^{1-\tau}$. The parameter λ_t controls the overall level of taxation (with a higher λ generating lower taxes) and τ controls the progressivity of the tax schedule. We allow the tax level λ_t to vary as a function of the government’s budgetary rule. We use $\tau = 0.166$ following Fleck et al. (2021)⁹ and estimate the steady-state λ to match a 40% average tax rate on an income of \$300,000.¹⁰

In addition to paying retirement benefits and collecting taxes, the government is in charge of collecting the assets of the dead and of endowing newborns with their initial assets. For aggregate accounting purposes, we denote the assets of households that die at time t with Λ_t and the total endowments to newborns, which are constant, with $\mathcal{E} \equiv \int k_{i,t} dD_t(\mathbf{a}_{i,t} = 26)$. The endowments of newborns are heterogeneous and match the distribution of wealth for 21 to 25 year olds in the 2019 SCF.

2.5 Preferences

Agents receive utility from consumption through a constant relative-risk aversion function,

$$u(c) = \frac{C^{1-\rho}}{1-\rho}.$$

Each period, they face a probability of death $\delta_{\mathbf{a}_{i,t}}$ that is taken from the SSA life tables. Upon death, agents receive utility from leaving their wealth as a bequest through the function

$$\phi(a) = b \times \frac{(a + \kappa)^{1-\rho}}{1-\rho},$$

where b controls the intensity of the bequest motive and κ the extent to which leaving bequests is a “luxury.” Finally, they receive disutility from their labor hours

$$v(n) = \varphi \times \frac{n^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}},$$

where ν is the Frisch elasticity of labor supply and φ is a scaling factor.

2.6 Timing and Recursive Formulation

Figure 2 summarizes the timing of the events that happen within a year from the point of view of the household. First, aggregate shocks occur and define the return that agents receive on their savings from the last period. Mortality is then realized, and a fraction of the agents dies and receives utility from their bequests, which the government then collects. Newborns enter the model with $\mathbf{a} = 26$, and the government endows them with assets that are random and follow a distribution calibrated to that of the SCF for respondents between the ages of 21 and 25. The agents then receive their wages, pension benefits, and transfers,

⁹We use their estimate for pooled U.S. data and a narrow definition of transfers (Table 2, bottom row).

¹⁰This target also comes from Fleck et al. (2021) in Figure 8 of the main text.

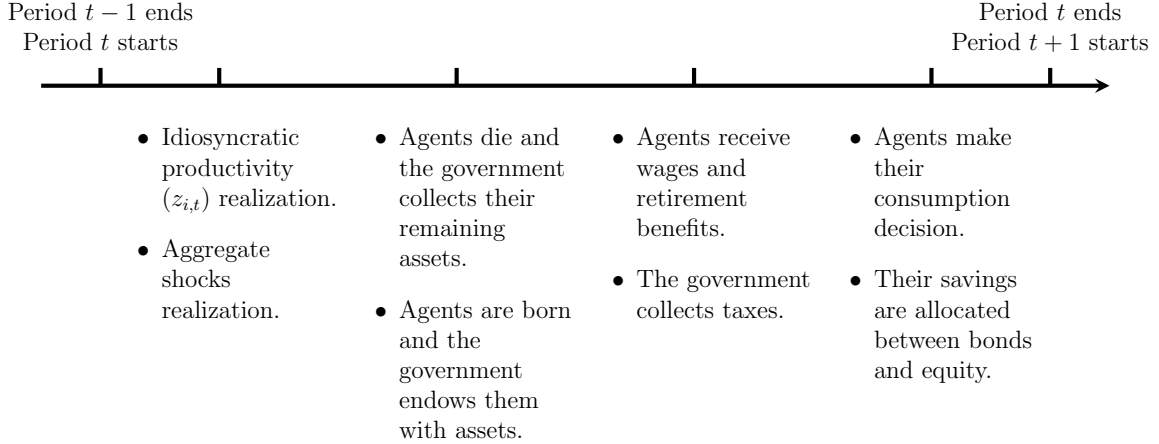


Figure 2: Summary of timing in the model

and the government collects taxes. Finally, the agents decide how much to consume, and their remaining assets are distributed between government bonds and equity in the firm following the rule described in Section 2.3.

We now specify the recursive formulation of the optimization problem of an individual that takes the sequence of wages, return factors, hours, and taxes $\{w_{t+s}, R_{t+s}^b, R_{t+s}^s, L_{t+s}, \lambda_{t+s}\}_{s=0}^{\infty}$ as given. We omit individual subscripts i . We start by defining $V_{a,t}(\alpha, z_t, a_{t-1})$, which is the value an agent expects at the beginning of the period before knowing whether he will survive or not. The value is

$$\begin{aligned}
 V_{a,t}(\alpha, z_t, a_{t-1}) &= \delta_{a_t} \phi(k_t) + \delta_{a_t} \tilde{V}_{a,t}(\alpha, z_t, k_t) \\
 \text{Where:} & \\
 \tilde{R}_t &= R_t^b + \zeta(\mathbf{a}_t, a_{t-1}) \times (R_t^s - R_t^b) \\
 k_t &= \tilde{R}_t \times a_{t-1},
 \end{aligned} \tag{7}$$

where k_t denotes assets after capital returns and $\tilde{V}_{a,t}(\alpha, z_t, k_t)$ is the value that the agent expects conditional on survival.

The definition of the value function conditional on survival changes with an agent's working status, as productivity z_t becomes deterministic upon retirement. For an agent that

is still working ($\mathbf{a} \leq 65$),

$$\begin{aligned} \tilde{V}_{\mathbf{a},t}(\alpha, z_t, k_t) &= \max_{c_t} u(c_t) - v(n_t) + \beta \mathbb{E}_t[V_{\mathbf{a}+1,t+1}(\alpha, z_{t+1}, a_t)] \\ \text{subject to} \\ m_t &= k_t + \mathcal{T}(w_t \times \tilde{y}_t \times n_t) \\ \tilde{y}_t &= \exp\{\alpha + f_{\mathbf{a}} + z_t\} \\ a_t &= m_t - c_t \\ a_t &\geq 0 \\ z_{t+1} &\sim \Pi_{\mathbf{a}_t}(z_t) \end{aligned} \tag{8}$$

and for an agent that has retired ($\mathbf{a} > 65$),


$$\begin{aligned} \tilde{V}_{\mathbf{a},t}(\alpha, z_t, k_t) &= \max_{c_t} u(c_t) + \beta \mathbb{E}_t[V_{\mathbf{a}+1,t+1}(\alpha, z_{t+1}, a_t)] \\ \text{subject to} \\ m_t &= k_t + \mathcal{T}(d_{i,t}) \\ d_t &= \exp\{\alpha + f_{\mathbf{a}} + z_t\} \\ a_t &= m_t - c_t \\ a_t &\geq 0 \\ z_{t+1} &= z_t. \end{aligned} \tag{9}$$

At the terminal age of 100, the survival probability becomes 0, and the agent’s value function is $V_{100,t}(\alpha, z_t, a_{t-1}) = \phi(\tilde{R}_t \times a_{t-1})$, where $\tilde{R}_t = R_t^b + \zeta(100, a_{t-1}) \times (R_t^s - R_t^b)$.

2.7 Calibration of the Household Model

We calibrate our household model to replicate various features of the distribution of income and wealth at different ages in the 2019 wave of the SCF. We use the SCF “summary files” and use `typewriter font` to denote variables defined in them. Our sample consists of respondents above the age of 21 that report a strictly positive income, and we use survey weights in all of our calculations.

We start by calibrating our income process, which requires a sequence of age-specific intercepts $\{f_{\mathbf{a}}\}_{\mathbf{a}=26}^{100}$ and a distribution of individual fixed-effects α . Our measure of income is the sum of wage and salary income (`wageinc`), and Social Security and pension income (`ssretinc`). The age-specific intercepts come from regressing the logarithm of income on a 5th-degree polynomial of age and predicting the fitted values for each age. For the distribution of individual fixed effects, we use a normal distribution discretized with three equiprobable points, $\alpha_i \sim \mathcal{N}(0, \sigma_\alpha)$. We estimate σ_α to match the dispersion of income across the life cycle. We form five-year age bins ($[26, 30]$, $[31, 35]$, ..., $[91, 95]$) and calculate the 25th, 50th, and 75th percentiles of income. Then, we find the σ_α that minimizes the distance between the model-implied percentiles and those in the data. Figure 3 displays the



Our measure of income is the sum of wages and salaries, and Social Security and pension income from the 2019 SCF. Our sample is households that report a strictly positive income and where the respondent is at least 21 years old. We sort households into the reported age bins and calculate the reported income percentiles for each age bin. Black crosses correspond to the model-fitted counterparts to these percentiles.

Figure 3: Age-Profiles of the Income Distribution in the Data and in the Model.

empirical and model-implied percentiles of the income distribution, showing that our model replicates the age patterns of both its level and dispersion.

Productivity dynamics, especially their persistence, can have a material effect on the behavior of consumption and savings. The optimal saving rate of a low-income household is very different depending on where it thinks its income will converge over time, at what rate, and with what degree of certainty. Figure A.2 illustrates the persistence of productivity (\tilde{y}) and shocks (z) in our model, tracing out the expected percentiles of each that a households expect to be in, given their percentile at age 30. The left panel of the figure shows that productivity is quite persistent: a household in the 20th percentile of the age 30 income distribution can expect to retire having reached only the 26th percentile of the age 65 distribution, and a household that starts in the 80th percentile expects to reach the 74th. This persistence is due in large part to the individual fixed effects in productivity (α). The right panel of Figure A.2 removes the fixed effect and looks only at the distribution of persistent shocks z . Households starting in the 20th and 80th percentiles of the shock at age 30 expect to reach age 65 in the 50th and 60th percentiles, respectively.

We calibrate the preference parameters $\{\rho, \beta, b, \kappa\}$ matching the age patterns in the distribution of savings. For the same five-year age bins used in calibrating the income process, we find the 25th, 50th, and 75th percentiles of the ratio of financial assets (fin) to our measure of income. We find the preference parameters that minimize the distance between

Table 1: Estimated Household Parameters

Parameter	Symbol	Estimate
Preferences		
Relative-Risk Aversion	ρ	1.88
Discount Factor	β	0.93
Bequest Intensity	b	299.69
Bequest Shifter	κ	11.03
Income Process		
Income F.E. Std. Dev.	σ_α	0.62

We estimate the preference parameters matching the age profiles of the 25th, 50th, and 75th percentiles of the wealth-to-income ratio. Wealth-to-income ratios in the SCF are the ratio of financial assets to our measure of income (wages, salaries, Social Security, and pension income). In our model, wealth-to-income ratios are end-of-period assets divided by income, $a_{i,t}/y_{i,t}$. We estimate the income process parameter independently to match the same percentiles of the age profiles of income.

the age-varying percentiles implied by the model and those in the SCF.¹¹¹² Our preference estimates, which we present in Table 1, lie in ranges that are typical for similar exercises in the labor economics and macroeconomics literatures: a coefficient of relative risk aversion close to 2 that implies an intertemporal elasticity of substitution close to 1/2, an annual discount factor above 0.9, and an intense bequest motive that increases in importance for wealthier agents ($\kappa > 0$).¹³ Figure 4 compares the empirical and model-implied percentiles of the wealth-to-income ratio, demonstrating that our parsimonious model achieves a remarkable fit of the age profiles of wealth and its dispersion. The main shortcoming of the model is its poor fit of wealth above the median before the age of 45: it prescribes that agents up to this age must hold minimal savings, and while most do, there are some who do not. We suspect that this shortcoming is linked to the absence of transitory income shocks in the current version of the model; these shocks are important for generating the “buffer-stock” motive, which is the main source of savings before the age of 45 in these models (see Carroll, 1997; Gourinchas & Parker, 2002).


2.8 Wealth Distribution and MPCs

The literature on heterogeneous-agent macroeconomics has shown that the distributions of savings and MPCs across agents in the economy modulate the effects of fiscal and monetary policy (Kaplan & Violante, 2014; Carroll et al., 2017; Auclert, 2019). Studies in this literature

¹¹We use the standard Simulated Method of Moments loss function, with a diagonal weighting matrix that roughly rescales all moments to have similar magnitudes.

¹²We use the distribution of end-of-period assets $a_{i,t}$ in our model as the counterpart of SCF financial assets for this exercise.

¹³See, for example, Carroll (1992), Attanasio et al. (1999), Gourinchas and Parker (2002), Cagetti (2003), and De Nardi et al. (2010).



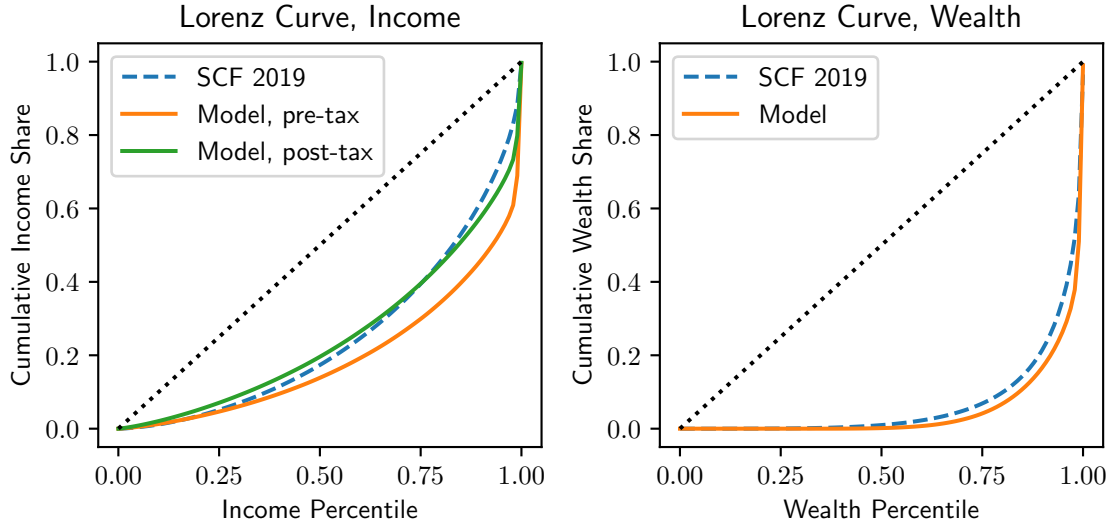
Wealth-to-income ratios in the SCF are the ratio of financial assets to our measure of income (wages, salaries, Social Security, and pension income). In our model, wealth-to-income ratios are end-of-period assets divided by income, $a_{i,t}/y_{i,t}$. We sort households into the reported age bins and calculate the reported percentiles of the wealth-to-income ratio for each age bin. Black crosses correspond to the model-fitted counterparts to these percentiles.

Figure 4: Age Profiles of the Savings Distribution in the Data and in the Model

and in household finance have also demonstrated that off-the-shelf models cannot reproduce various important features of these distributions. This section shows that our life-cycle model of households can reproduce various features of the wealth and MPC distributions that have been deemed important and difficult to reproduce.

Wealth distribution. The first non-targeted feature of interest that our model reproduces is the degree of inequality of the wealth distribution, which is much greater than that of the income distribution. Figure 5 depicts Lorenz curves for income and wealth that aggregate individuals of all ages, comparing the model-implied curves with those of the SCF. Table 2 reports the Gini coefficients associated with each of the curves. Our calibrated process generates a distribution of income with a similar degree of inequality to that of the SCF. The progressive tax system discussed in Section 2.4 makes the after-tax distribution of income substantially less unequal than the pretax distribution. Traditional models that feature agents with identical preferences and rely on precautionary savings often struggle to replicate the empirical reality of a wealth distribution that is much more unequal than the income distribution (Quadrini & Rios-Rull, 1997; Stachurski & Toda, 2019).¹⁴ Figure 5 and Table 2 demonstrate that our model fits this feature of the wealth distribution, replicating its

¹⁴Various studies have allowed features of the earnings process (Castañeda et al., 2003) or specifications of heterogeneous preferences (Carroll et al., 2017) to vary in a way that they estimate to match wealth inequality.



The model statistics reported in this figure come from the steady state of our model, aggregating over households of all ages. The black dotted lines represent perfect equality. Our measure of income in the SCF is the sum of wages and salaries, and Social Security and pension income. Our measure of wealth is total financial assets. The model counterpart of wealth is end-of-period assets, $a_{i,t}$.

Figure 5: Lorenz Curves for Income and Wealth in the SCF Data and in our Model.

Lorenz curve and associated Gini coefficient with great fidelity. The success of our model in this dimension comes from explicitly modeling the life cycle of agents and from the luxury bequest specification, both of which are known to improve the predictions of this class of models about the distribution of wealth (Huggett, 1996; De Nardi, 2004).

Figure 6 provides further insights into the roles of life-cycle and precautionary motives for saving in our model. The black dotted lines show average saving (total income minus consumption) and wealth by age in our benchmark model. Saving increases steadily with age and peaks before retirement. Saving declines through retirement but remains positive except for the oldest households, implying that average wealth peaks well into retirement. We illustrate the relative importance of different saving motives following Gourinchas and Parker (2002). The blue lines labelled “life cycle” correspond to a version of the model without income risk, where every agent receives the age-specific average of post-tax income that individuals with his productivity fixed-effect α_i would receive in the baseline model. This “life cycle” version retains every other feature of the model, including bequest motives and the borrowing constraint. The orange lines labelled “buffer” capture the precautionary motive which we compute as the difference between the baseline values and the life cycle values. As in Gourinchas and Parker (2002), young households save exclusively out of precaution, saving for retirement starts around the age of 40, and life-cycle motives become the dominant reason for saving after 55. All in all, life-cycle motives account for the bigger share of aggregate wealth.

Table 2: Wealth and Income Inequality

Variable	Gini Coefficient	
	Data (SCF 2019)	Model
Pretax Income	0.50	0.61
Post-tax Income	-	0.50
Wealth	0.86	0.90

Gini coefficients are the areas between the equality diagonal and the Lorenz curves in Figure 5. Our measure of income in the SCF is the sum of wages and salaries, and Social Security and pension income. Our measure of wealth is total financial assets. The model counterpart of wealth is end-of-period assets, $a_{i,t}$. The model statistics reported in this table come from the steady state of our model, aggregating over households of all ages.

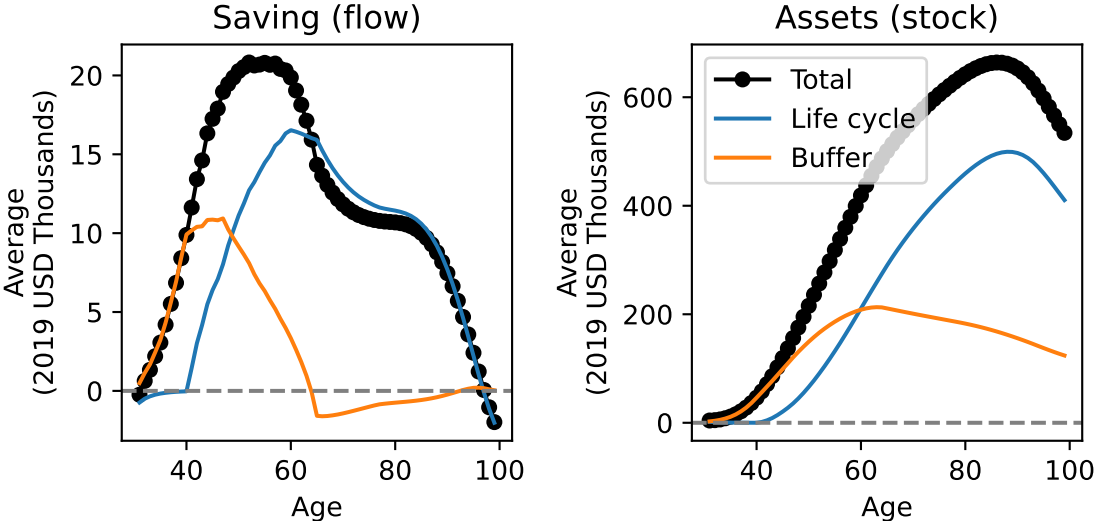
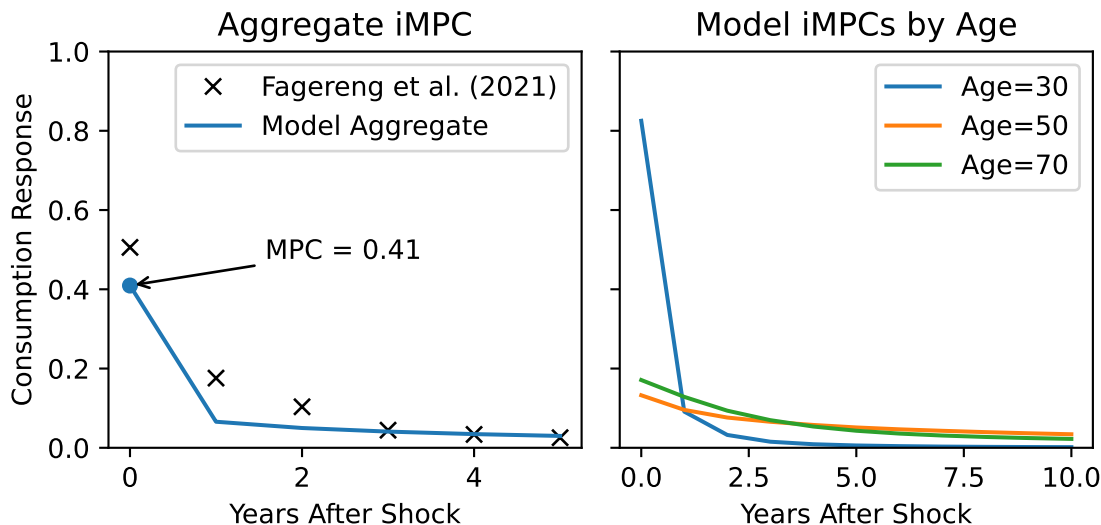


Figure 6: The Role of Different Mechanisms for Wealth Accumulation



The figure presents intertemporal marginal propensities to consume, which are the consumption responses at different horizons to a lump-sum transfer received at time 0. The left panel compares the aggregate iMPC of the household sector in our model to the estimates of Fagereng et al. (2021). The right panel disaggregates iMPCs in our model by the age of the agents at the time of the transfer (time 0). For every line, the value at time 0 is the marginal propensity to consume.

Figure 7: Intertemporal Marginal Propensities to Consume (iMPCs).

Average MPCs. The second feature that we highlight is that our model generates MPCs of magnitudes that are consistent with empirical estimates. The growing empirical literature that measures households’ MPCs has obtained estimates that, for the average annual MPC, range roughly between 0.2 and 0.6;¹⁵ this is an order of magnitude greater than the MPCs predicted by standard representative-agent models. The left panel of Figure 7 depicts the aggregate *intertemporal* MPCs (iMPCs) of the household sector in our model, comparing it with the empirical estimates of Fagereng et al. (2021). The MPC implied by our model, which is the $t = 0$ value of the iMPCs, falls squarely within the 0.2 to 0.6 range of empirical estimates and close to the point estimate of Fagereng et al. (2023). This is a notable feat given that our model is calibrated to match the age profiles of a broad measure of savings (total financial assets as a multiple of income, see Figure 4). Indeed, in their review of heterogeneous-agent models and MPCs, Kaplan and Violante (2022) conclude that one-asset models where agents save mainly out of precaution can only generate MPCs as high as their empirical counterparts if they are calibrated to match narrow measures of liquid wealth. Our model sidesteps this trade-off because we include additional reasons for saving: life-cycle fluctuations in earnings, and bequests.

MPC heterogeneity. A third desirable feature of our model is that it generates substantial heterogeneity in MPCs without relying on heterogeneous or behavioral preferences;

¹⁵See Jappelli and Pistaferri (2010), Carroll et al. (2017), and Crawley and Theloudis (2024) for summaries of this literature and Fagereng et al. (2021) for estimates using Norwegian administrative data.

instead, the variation is mainly driven by wealth and age. Figure 8 depicts the average annual MPC for agents of different ages and in different quintiles of the age-specific distribution of wealth. The figure shows that in our steady-state distribution there are agents with MPCs lower than 0.05 and higher than 0.95, in spite of having identical preferences. The group of agents with no savings and MPCs close to 1.0—the “hand-to-mouth”—are important for generating the high aggregate MPC of our model. In steady state, 34% of the agents in our model are hand-to-mouth, which is close to the 40% empirical estimates of Aguiar et al. (2020) and McKay and Wolf (2023).

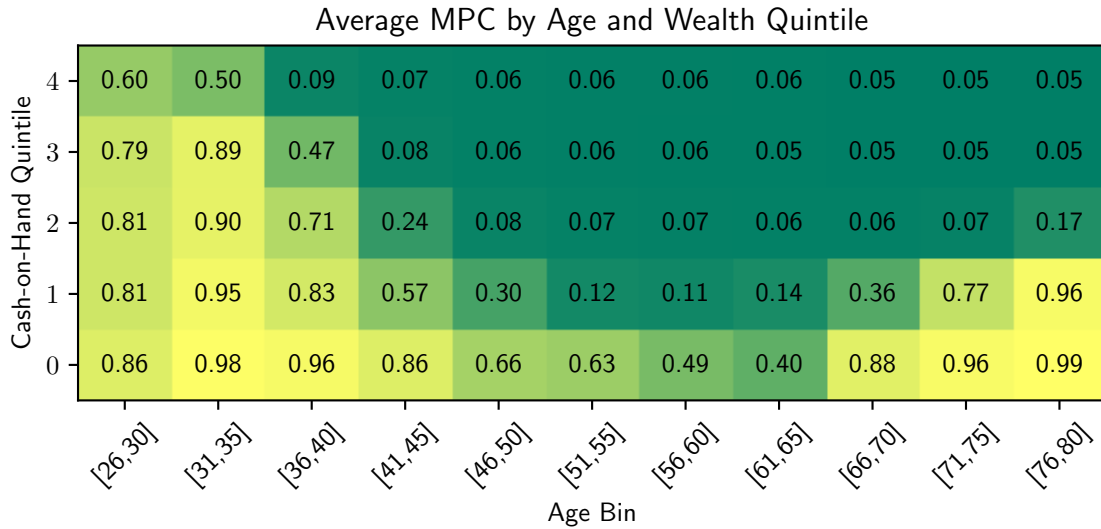
The heterogeneous MPCs in Figure 8 co-vary with age and wealth: they fall sharply with age up to retirement and also with wealth for any given age group. These patterns qualitatively match the conclusion from Fagereng et al. (2021) that age and liquid assets are the main household characteristics that systematically correlate with households’ MPCs. After retirement, the luxury-bequest motive prevents a fraction of the agents from running down their assets: only around half of agents do. This prevents the average MPCs of the old from increasing as sharply as they do in, for example, Carroll et al. (2017), Braun and Ikeda (2021), and Bullard et al. (2023), and also replicates the fact that a fraction of the elderly keep large stocks of savings (De Nardi et al., 2010).¹⁶

Figure 9 highlights the importance of income inequality *within age* and bequests motives in generating an empirically realistic age-profile of MPCs. The blue line labelled “no inequality” corresponds to a version of the model in which each household receives the average post-tax income at every age. The orange line labelled “no bequests” comes from a version of the model that shuts down the bequest motive by setting $b \rightarrow 0$ and $\kappa \rightarrow \infty$. Finally, the green line labelled “no inequality or bequests” implements both of these changes together. We re-estimate the preference parameters in these specifications targeting the age profiles of wealth-to-income ratios, as we do for the baseline model. The only difference is that, in models that do not feature inequality, we target only the median wealth ratios of the relevant age bins. The fit of the baseline model is showcased in Figure 4 and that of alternative specifications in Figure C.2 of Appendix C.

In our baseline model, MPCs fall gradually by age, and rise only moderately in old age. Without income and wealth inequality within age groups, young households quickly run down their inheritance and become hand-to-mouth. The pressure to save for retirement reaches a critical point around the age of 42. As a result, the MPC drops sharply and stays low until retirement. Bequests in turn, are crucial to prevent the MPC of old households from rising sharply towards one at the end of life.

Taking stock. Our model can generate substantial heterogeneity in wealth and MPCs across and within age groups that aligns well with the available evidence. This makes it well-suited for studying the redistribution channel of monetary (and fiscal) policy. Auclert (2019) demonstrates that the redistribution channel of the effect of monetary policy on ag-

¹⁶Bequests and medical expenditures are the two main reasons that have been postulated for the high saving rates of *some* of the elderly (see, for example, De Nardi et al., 2010; Ameriks et al., 2020). Adding medical expenditure risks to our model would further depress the MPCs of the old.



The figure depicts annual MPCs across the steady-state distribution of households in our model. Starting with the steady-state distribution, we group agents into the depicted age bins. For each age bin, we group agents into quintiles of their cash-on-hand ($m_{i,t}$). For each age and cash-on-hand group, we find the average MPC across agents, measuring the MPC as the derivative of their consumption function, $\partial c_{i,k} / \partial k_{i,t}$.

Figure 8: MPCs Across the Life Cycle and the Wealth Distribution.

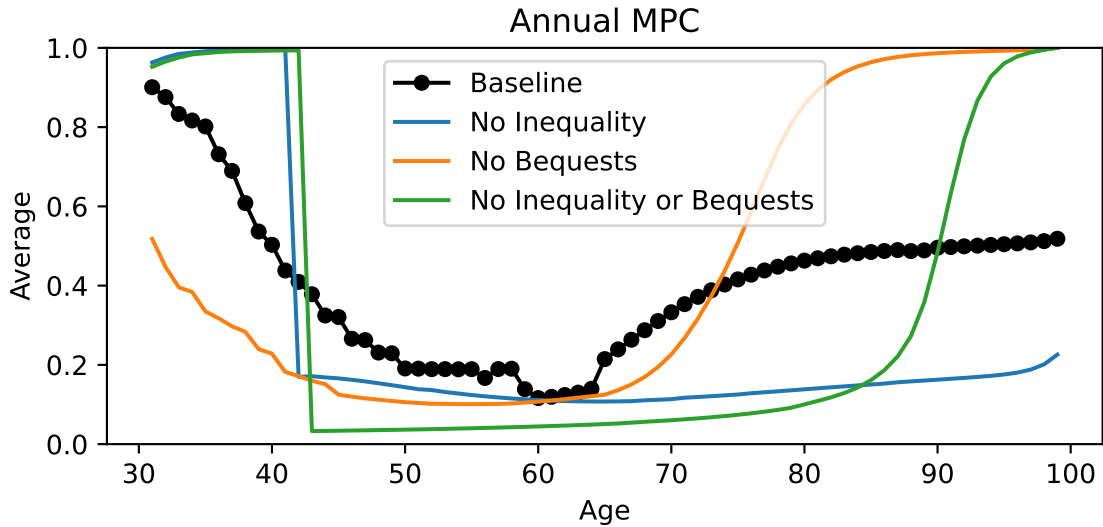


Figure 9: MPCs in Alternative Models that Target Life-Cycle Wealth Ratios

gregate consumption operates through the covariance of agents' MPCs with various features of their income, consumption plans, and asset holdings.¹⁷ However, previous studies of monetary policy in heterogeneous-agent economies have abstracted away either from life-cycle considerations or from heterogeneity within age groups.¹⁸ When models incorporate *only* life-cycle heterogeneity, every household in a cohort owns the average wealth of the cohort and this makes their MPCs much lower than the empirical evidence suggests (see, for example, Braun & Ikeda, 2021). Instead, our model features a rich representation of agents' life cycles and, at any given age, heterogeneity in their savings and portfolios. This heterogeneity and the high average MPCs of our model make it a good setting for studying the redistributive effects of monetary policy and their consequences for aggregate demand and output.

3 A Life-Cycle HANK Model

Next, we embed our model of overlapping generations of heterogeneous households into general equilibrium. Our goal is to study the transmission mechanism and distributive effects of monetary policy in the presence of uninsurable income risk and life-cycle considerations. To maximize comparability with the existing HANK literature, the rest of our model stays close to Auclert et al. (2018) and Alves et al. (2020).¹⁹

3.1 Setup

Time is discrete, and each period corresponds to a year. The economy consists of a unit mass of heterogeneous households, a unit mass of labor unions, a unit mass of firms, a central bank, and a government.

Households. The household decision problem is described in detail in section 2. From a macroeconomic perspective, the household sector is a mapping from aggregate sequences $\{w_t, L_t, R_t^b, R_t^s, \lambda_t\}_{t=0}^{\infty}$ to aggregate sequences $\{C_t, A_t, T_t, \Lambda_t, v'(N_t^*), u'(C_t^*)\}_{t=0}^{\infty}$. The inputs are the real wage w_t , labor demand L_t , return on nominal bonds R_t^b , return on stocks R_t^s , and the intercept of the retention function λ_t . The outputs are consumption C_t , asset demand A_t , taxes net of transfers T_t , bequests Λ_t , marginal cost of labor $v'(N_t^*)$, marginal benefit of labor $u'(C_t^*)$. Let D_t denote the measure of households over idiosyncratic states $(\mathbf{a}, \alpha, z, k)$ in period t .²⁰ Then, aggregate outputs are defined as follows, where we use i subscript

¹⁷Specifically, their incomes, net nominal positions, and unhedged interest rate exposures.

¹⁸See Kaplan et al. (2018) and Auclert et al. (2020) for examples without a life-cycle component, and Braun and Ikeda (2021) and Bielecki et al. (2022) for examples with life-cycle differences but homogeneous cohorts.

¹⁹The model of Alves et al. (2020) is a quantitative refinement of the seminal work of Kaplan et al. (2018) that allows for unequal exposure to aggregate income fluctuations and capital adjustment costs.

²⁰Recall that \mathbf{a} is age, α is income fixed effect, z is persistent income shock, and k is wealth after revaluation due to unanticipated shocks.

to highlight objects that vary across individuals without writing out their dependence on idiosyncratic states:

$$C_t = \int c_{i,t} dD_t, \quad A_t = \int a_{i,t} dD_t, \quad \Lambda_t = \int \delta_{i,t} k_{i,t} dD_t, \quad (10)$$

$$T_t = \int y_{i,t} - \mathcal{T}(y_{i,t}) dD_t, \quad y_{i,t} = \begin{cases} w_t \times \tilde{y}_{i,t} \times n_{i,t} & \text{for } \mathbf{a}_{i,t} \leq 65, \\ d_{i,t} & \text{for } \mathbf{a}_{i,t} > 65. \end{cases} \quad (11)$$

We define the marginal utilities $v'(N_t^*)$, $u'(C_t^*)$ below in the context of the wage Phillips curve.

Labor unions. Nominal wages are set by labor unions whose objective is to maximize the average value of working-age households. They take as given the consumption-saving decisions of individual households as well as the age-specific labor demand schedule $\gamma(\mathbf{a}, L)$. We model wage stickiness via a quadratic adjustment cost specified in utils, in the tradition of Rotemberg (1982). In appendix B.2, we describe the decision problem of unions in detail and show that it implies a wage Phillips curve

$$\pi_t^w (\pi_t^w - 1) = \kappa_w \left(\frac{v'(N_t^*)}{u'(C_t^*)} - 1 \right) + \beta \mathbb{E}_t [\pi_{t+1}^w (\pi_{t+1}^w - 1)], \quad (12)$$

where $\pi_t^w = \pi_t w_t / w_{t-1}$ is nominal wage inflation, $\kappa_w > 0$ is the slope of the wage Phillips curve, and $v'(N_t^*)/u'(C_t^*)$ is the average marginal rate of substitution between consumption and labor among working-age households. This is a sufficient statistic that captures the impact of the distribution and distortionary taxes on labor supply. Its components are defined as

$$v'(N_t^*) = \int L_t v'(n_{i,t}) \epsilon_w \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} dD_t(\mathbf{a}_{i,t} \leq 65), \quad (13)$$

$$u'(C_t^*) = \int u'(c_{i,t}) T'(y_{i,t}) y_{i,t} \left(\epsilon_w \frac{L_t}{n_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} - 1 \right) dD_t(\mathbf{a}_{i,t} \leq 65). \quad (14)$$

Firms. The production block has two types of firms. A competitive final goods firm aggregates intermediate goods with constant elasticity of substitution $\epsilon_p > 1$. Intermediate goods are produced by a unit mass of monopolistically competitive firms. These firms are identical ex ante. They have a Cobb-Douglas production function $y_t = F(k_{t-1}, l_t) = \Theta k_{t-1}^\alpha l_t^{1-\alpha}$, where Θ is TFP. We assume quadratic adjustment costs on both capital and the intermediate goods price.

In appendix B.1, we describe the decision problems of firms in detail and show that they give rise to a symmetric equilibrium in which all firms set the same price and produce with the same amount of labor and capital. The resulting inflation dynamics is characterized by a Phillips curve

$$\pi_t (\pi_t - 1) = \kappa_p (\mu_p m c_t - 1) + \mathbb{E}_t \left[\frac{\pi_{t+1} (\pi_{t+1} - 1) Y_{t+1}}{R_t^e Y_t} \right], \quad (15)$$

where $\kappa_p > 0$ is the slope of the Phillips curve, $\mu_p = \epsilon_p/(\epsilon_p - 1)$ is the desired markup of intermediate goods producers, $mc_t = w_t/F_L(K_{t-1}, L_t)$ is their real marginal cost, R_t^e is the gross interest rate between period t and $t + 1$, and Y_t is aggregate output.

The dynamics of investment is given by

$$Q_t = 1 + \psi \left(\frac{K_t}{K_{t-1}} - 1 \right), \quad (16)$$

$$R_t^e Q_t = \mathbb{E}_t \left[\alpha \frac{Y_{t+1}}{K_t} mc_{t+1} - \frac{I_{t+1}}{K_t} - \frac{\psi}{2} \left(\frac{K_{t+1}}{K_t} - 1 \right)^2 + \frac{K_{t+1}}{K_t} Q_{t+1} \right], \quad (17)$$

where Q_t is marginal Q, $\psi > 0$ is the capital adjustment cost, and $I_t = K_t - (1 - \delta)K_{t-1}$ is investment.

Intermediate goods firms make a positive profit in equilibrium, on account of their accumulated capital stock and monopoly power. The flow profit²¹ is

$$d_t = Y_t - w_t L_t - I_t - \frac{\psi}{2} \left(\frac{K_t}{K_{t-1}} - 1 \right)^2 K_{t-1}. \quad (18)$$

Firms are owned in part by households and in part by the government. We assume that the government receives a constant share of profits $\omega_g d_t$ as dividends, while households trade freely the remaining $1 - \omega_g$ share of equity. Absent portfolio adjustment costs and aggregate uncertainty, stocks and bonds must offer the same expected return, implying that

$$p_t = \mathbb{E}_t \left[\frac{d_{t+1} + p_{t+1}}{R_t^e} \right]. \quad (19)$$

Fiscal policy. The government pays pensions to retirees, provides endowments \mathcal{E} to newborn households, and consumes an exogenous amount G_t of the final good. It finances these expenditures by issuing one-period nominal bonds B_t , collecting bequests Λ_t from households that die, from dividends from its ownership of firms, and from running a progressive tax and transfer system. In sum, the primary surplus of the government is

$$S_t = T_t - G_t + \Lambda_t - \mathcal{E} + \omega_g d_t, \quad (20)$$

and its budget constraint is

$$B_t + S_t = R_t^b B_{t-1}. \quad (21)$$

We assume that the government adjusts the intercept of the retention function, λ_t , to implement the following path of the primary surplus:

$$S_t - \bar{S} = (1 - \rho_R) \left[(R_t^b - 1)B_{t-1} - (\bar{R}^b - 1)\bar{B}_{t-1} \right] + \rho_B (B_{t-1} - \bar{B}) + \rho_T (T_t - \bar{T}) + \rho_d \omega_g (d_t - \bar{d}), \quad (22)$$

where the coefficients ρ_\bullet govern the reaction to the primary surplus to fluctuations in different categories of revenues and expenditures.

²¹We don't keep track of the resource cost of price adjustments which is of second order and thus irrelevant given our solution method.

Monetary policy. The central bank sets the nominal interest rate according to the rule

$$R_t^n = R_{ss}^n + \phi_\pi(\pi_t - \pi_{ss}) + \varepsilon_t^{mp}, \quad (23)$$

where $\phi_\pi > 1$, and ε_t^{mp} is an exogenous monetary policy shock. The realized return on nominal bonds is thus

$$R_t^b = \frac{R_{t-1}^n}{\pi_t}. \quad (24)$$

Equilibrium. Given a sequence of monetary policy shocks $\{\varepsilon_t^{mp}\}$, an exogenous distribution of endowments, and initial conditions D_{-1}, K_{-1}, B_{-1} , equilibrium is a sequence of allocations $\{Y_t, C_t, A_t, B_t, \Lambda_t, K_t, I_t, L_t, d_t\}$ and prices $\{R_t^e, R_t^b, R_t^s, R_t^n, Q_t, w_t, mc_t, p_t\}$ such that

- Households, labor unions, and firms optimize;
- The government and the central bank follow their policy rules;
- Goods and asset markets clear

$$Y_t = C_t + G_{ss} + I_t + \frac{\psi}{2} \left(\frac{K_t}{K_{t-1}} - 1 \right)^2 K_{t-1}, \quad (25)$$

$$A_t = B_t + (1 - \omega_g)p_t. \quad (26)$$

3.2 Calibration

For the calibration of the household block, see sections 2.7 and 2.8. The non-household blocks of the model are calibrated to ensure that the household block is consistent with general equilibrium. Table 3 summarizes the calibration of macroeconomic aggregates and selected parameters. Next, we discuss this calibration block by block.

Labor unions. We calibrate disutility of labor, φ , to ensure that $v'(N^*) = u'(C^*)$ holds given the stationary distribution of households. We calibrate the slope of the wage Phillips curve based on the equivalent Calvo model

$$\kappa_w = \frac{1}{1 + \Gamma_w} \frac{[1 - \beta(1 - \xi_w)]\xi_w}{1 - \xi_w}, \quad (27)$$

where β is the discount factor inherited from the household block, $\xi_w = 0.33$ is the annual frequency of wage adjustment from Grigsby et al. (2021), and Γ_w is a real rigidity parameter that we set to 5 following Auclert et al. (2018). The elasticity of substitution is set to $\epsilon_w \rightarrow \infty$.

Table 3: Calibration of Macroeconomic Aggregates

Quantity	Symbol	Value
Macro Aggregates (Steady State)		
Return Factors	$\{R^e, R^b, R^s\}$	1.02
Consumption Ratio	C/Y	0.48
Wealth Ratio	A/Y	1.77
Capital Ratio	K/Y	2.16
Investment Ratio	I/Y	0.30
Dividends Ratio	d/Y	0.12
Government Spending Ratio	G/Y	0.23
Government Debt Ratio	B/Y	0.29
Equity Share of HH Assets	Ξ	0.84
Parameters		
Frisch Elasticity	ν	0.50
Taylor Rule Coeff. on Inflation	ϕ_π	1.50
Wage P.C. Slope	κ_w	0.03
Goods P.C. Slope	κ_p	0.24
Capital Share	α	0.34
Capital Adjustment Cost	ψ	0.86
Capital Dep. Rate	δ_k	0.14
Govt. Share of Firm	ω_g	0.31

See the main text for the targets and rationale behind the calibration. For macro aggregates, the table reports steady state-values.

Firms. The household block implies a real wage and total hours worked in efficiency units. We calibrate TFP, Θ , and depreciation rate, δ_k , to justify these values given a labor share of $1 - \alpha = 0.66$ and a capital-to-output ratio of $K/Y = 2.16$. These are conventional choices. We calibrate the investment adjustment cost, ψ , to target a partial-equilibrium semi-elasticity of investment $d \log(I_t)/dR_t^e = -5.00$, in line with the findings of Koby and Wolf (2020), and He et al. (2022). Turning to price setting, we normalize gross inflation to $\pi = 1$ and calibrate the slope of the price Phillips curve based on the equivalent Calvo model, taking the annual frequency of price adjustment, $\xi_p = 0.67$, from Nakamura and Steinsson (2008). Analogously to wages, we allow for a real rigidity parameter $\Gamma_p = 5$. For market power, we take the limit $\mu_p \rightarrow \infty$ to shut down profits from monopoly power. Given this choice, equity price is equal to the value of capital, $p = K$.

Policy. The household block pins down total assets, A , as well as its share in stocks Ξ . The firm block pins down the total value of stocks, p . Together, these pin down the government's share of equity, ω_g . We calibrate government spending G , and debt B to be consistent with the primary surplus of the government. As our baseline, we assume that the government runs a balanced budget, setting $\rho_R = \rho_B = \rho_T = \rho_d = 0$. Turning to monetary policy, we set $\phi_\pi = 1.5$, a conventional value. In our main experiments, we assume that the monetary policy shock follows an AR(1) process with an annual autocorrelation of 0.5.

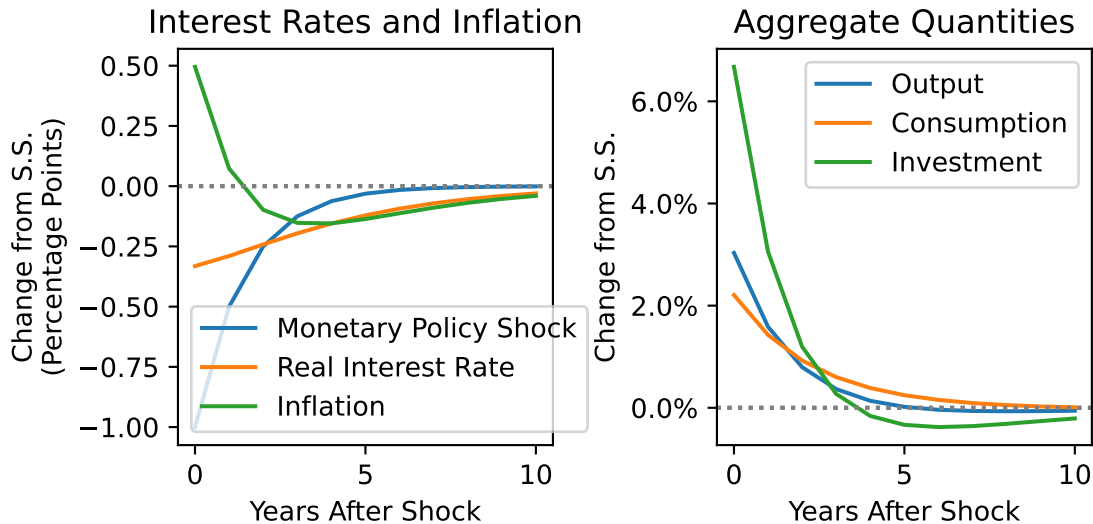
4 Monetary Policy in a Life-Cycle HANK

We study the response of the economy to an expansionary monetary-policy shock. The economy starts in steady state. At $t = 0$, a negative 1-percentage point shock to the Taylor rule (Equation 23) is announced. The shock decays at a 50% annual rate, $\{\varepsilon_t^{mp}\}_{t \geq 0} = \{-0.01 \times 0.5^t\}_{t \geq 0}$. We use the sequence-space Jacobian method of Auclert et al. (2021) to calculate the general-equilibrium responses of macroeconomic aggregates to the shock.

4.1 Aggregate Responses and Transmission Mechanisms

We start by describing the effects of the monetary expansion on macroeconomic aggregates and quantifying the importance of different channels for the response of consumption. The expansionary shock generates inflation and a persistent fall in real rates of asset returns, after an initial appreciation of real assets. Consumption, output, and investment increase, with the magnitude of their initial responses following that order. As other studies in the HANK literature have found, the main mechanisms behind the consumption response are indirect income effects arising from buoyant labor markets and changes in taxation. The effect of intertemporal substitution contributes a much smaller share of the consumption response.

The response of interest rates, inflation, and macroeconomic aggregates to the expansionary shock are qualitatively similar to those found in the New Keynesian literature (for



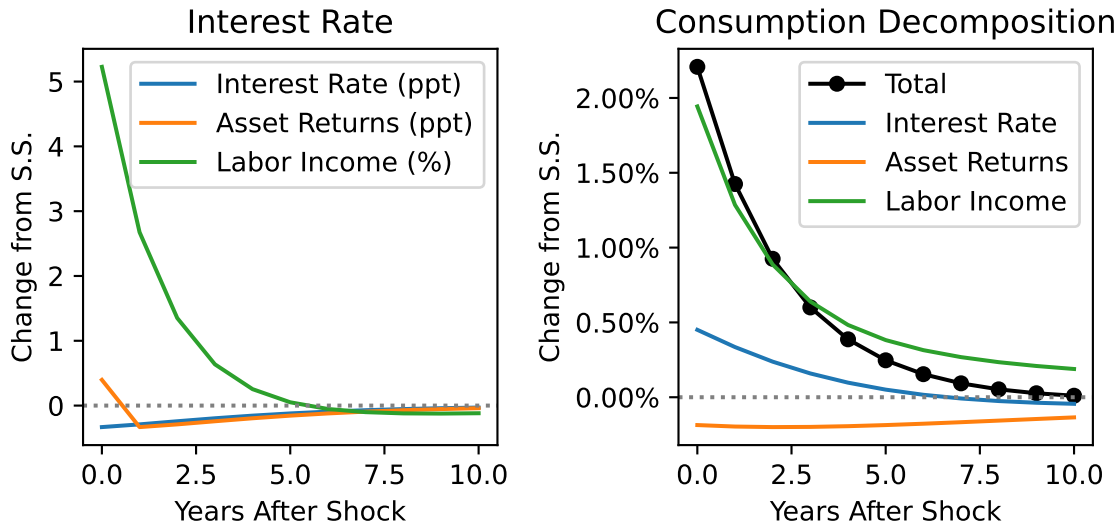
“Real Interest Rate” is R_t^e , the ex-ante rate used to discount future payments. “Inflation” is goods inflation.

Figure 10: Aggregate Responses to an Expansionary Monetary Policy Shock

example, Christiano et al., 2005; Kaplan et al., 2018).²² The left panel of Figure 10 depicts the path of the monetary policy shock (ε^{mp}) alongside the responses of the ex-ante interest rate (R^e) and goods inflation (π). Inflation jumps by 50 basis points and then decreases progressively, reaching its trough at 15 basis points below its steady-state value in years 3 and 4 after the shock, and then converging back slowly. The ex-ante real rate falls by 33 basis points on impact and remains below its steady-state level. The right panel of Figure 10 shows the response of output, consumption, and investment, which all increase on impact. The response of investment is the greatest, at 6.7%, followed by output at 3.0% and consumption at 2.2%.

We now examine the drivers of the consumption response to the expansionary shock. We dissect the response into three components. First, the intertemporal substitution response to changes in *expected* interest rates, holding the *realized* rates of return on households’ assets constant. Second, the response to changes in the realized rates of return on households’ assets, holding their expected rates constant. And, third, changes in households’ labor income net of taxes and transfers (their non-capital income). The left panel of Figure 11 depicts the trajectories of the ex-ante real rate, the realized rate of return on the total assets held by households, and total labor income (wages plus transfers net of taxes) paid to the household sector. The realized returns to total assets initially increase due to the repricing of stock holdings; after the initial shock, realized asset returns are equal to the lagged ex-ante interest rate. The figure shows that the monetary policy shock produces a

²²Just like Kaplan et al. (2018), our model does not generate the hump-shaped responses that empirical studies typically find. “Sticky expectations” (Carroll et al., 2020) are a promising way to generate hump-shaped responses in HANK models; see Auclert et al. (2020).



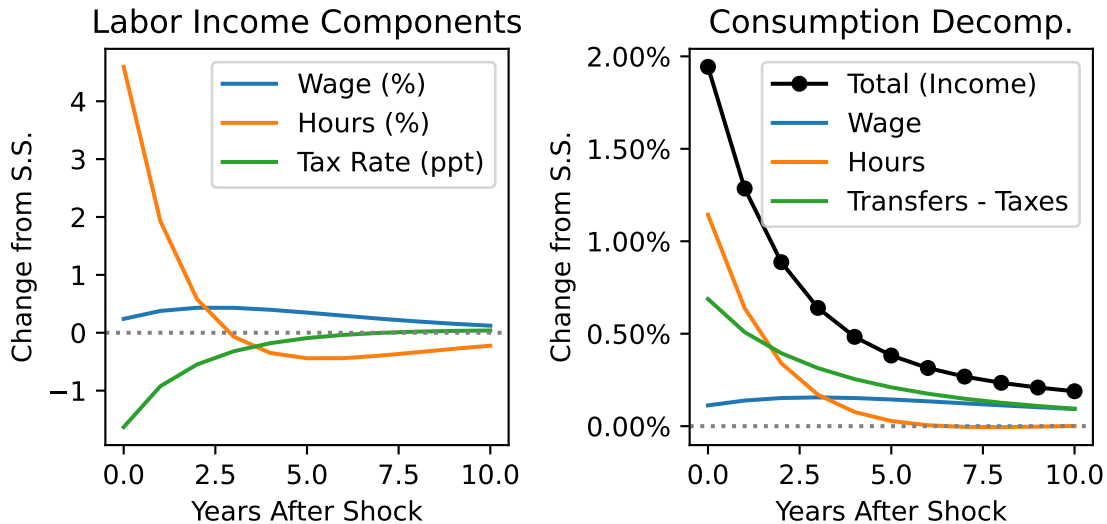
“Interest Rate” denotes the ex-ante real rate, R^e . “Asset Returns” denotes the realized returns on total assets held by households. “Labor Income” is the total income net of taxes and transfers paid to the household sector, $w \times L + T$. The left panel presents the response of each of these variables. The right panel presents the response of aggregate consumption when the deviations in these variables are given to the household sector independently (leaving all other values in steady state) and together (which produces the “Total” response).

Figure 11: Decomposition of Aggregate Consumption Response to a Monetary Policy Shock

5.2% increase in the labor income of households. Since our model features high MPCs (see Section 2.8), the changes in income and realized asset returns can have a much greater effect on consumption than in representative-agent models, where intertemporal substitution is the primary mechanism. We now evaluate this proposition.

Most of the increase in consumption that follows the monetary policy shock is due to the increase in non-capital income that households receive; intertemporal substitution has a much smaller positive effect, and reduced realized asset returns have a small negative effect. The right panel of Figure 11 decomposes the total response of consumption into the parts that can be adjudicated to each of these three mechanisms.²³ The figure shows that, by itself, the labor income channel produces a 1.9 percentage point initial increase in consumption out of the 2.2% total, and that it continues to be the main driver of the total response as time passes; this is consistent with the findings of, for example, Ampudia et al. (2018). By contrast, intertemporal substitution generates only a 0.5 percentage point increase on impact: less than one-quarter of the total response. The depressed realized returns on households’ assets generate a small decrease in consumption that starts at -0.2 percentage points and converges to 0 very slowly. Thus, in our model just as in, for example, Kaplan et al. (2018), Alves et al. (2020), and Auclert et al. (2020), most of the consumption response to the

²³In this and the following decompositions, we isolate the effect of a set of variables by calculating the response of the household sector to the perturbed sequence of that set of variables only, leaving all the other aggregate variables fixed in their steady state values.



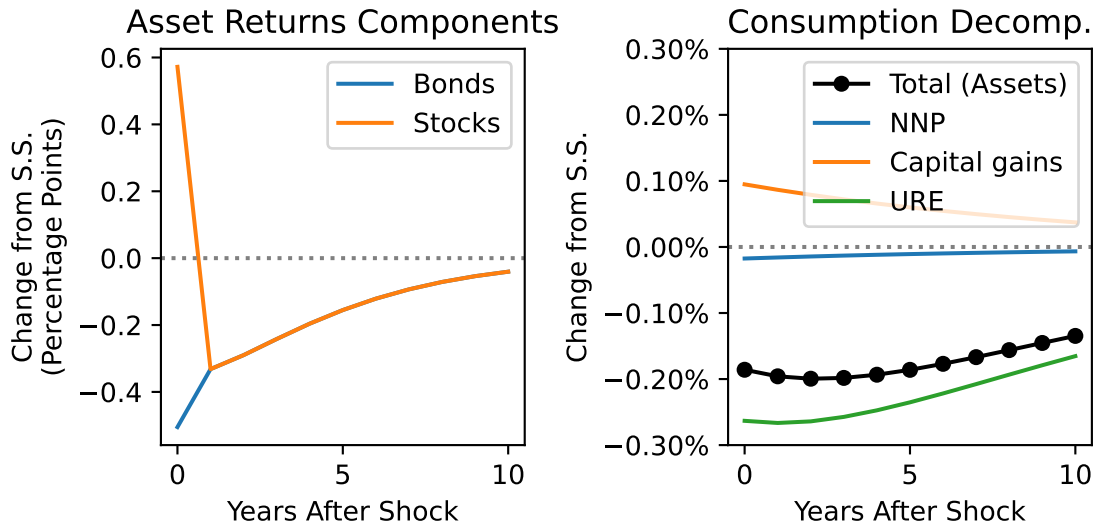
“Wage” denotes the real wage w . “Hours” denotes aggregate labor hours L . “Tax Rate” is $1 - \lambda$ from Equation 6. The left panel presents the response of each of these variables. The right panel presents the response of aggregate consumption when the deviations in these variables are given to the household sector independently (leaving all other values in steady state) and together (which produces the “Total” response).

Figure 12: Decomposition of Aggregate Consumption Response to a Monetary Policy Shock: Labor Income

monetary policy shock is due to indirect general-equilibrium effects that alter the disposable income of households.

Further decomposing the consumption response, we find that more than half of the income channel (and the total response in turn) is due to a buoyant labor market that increases employment hours and wages, and that a smaller but still significant share is due to tax reductions. The left panel of Figure 12 disaggregates the components of household labor income into hours, wages, and tax rates. Increases in the demand for labor generate initial jumps of 4.6% in total effective hours of work L and of 0.2% in the wage rate. These two mechanisms account for 1.3 percentage points out of the total 2.2% increase in consumption. With lower real interest payments and a larger taxable base, the budget rule (Equation 22) pushes the government to lower taxes: our measure of the tax rate, $1 - \lambda$, falls by 1.6 percentage points, and this contributes a non-negligible 0.7 percentage point to the initial consumption response. Because households are non-Ricardian and many have high MPCs, fiscal adjustments—and the assumed rules that they follow—become important contributors to the transmission of monetary policy (Kaplan et al., 2018; Alves et al., 2020).

The small negative effect of realized asset returns on consumption comes mainly from the unhedged interest rate exposure (URE) channel which is partly offset by the initial revaluation of stocks. Figure 13 decomposes this channel: the left panel shows that after the initial revaluation of stocks, R^b and R^s equalize. In the right panel, the net nominal position (“NNP”) and “capital gains” lines isolate the effects of an R_t^b and an R_t^s that turn



The left panel presents the trajectory of the realized returns of bonds R^b and stocks R^s . The right panel presents the response of aggregate consumption when different components of the asset returns channel are given to the household sector independently (leaving all other values in steady state) and together (which produces the “Total” response). “NNP” stands for net nominal position and represents a scenario where only the initial return in bond returns, R_t^b , is passed to the household block; all other variables and $\{R_{t+s}^b\}_{s \geq 1}$ are left in their steady-state values. “Capital Gains” isolates the effect of initial equity revaluation, only the initial return to stocks R_t^s changes, and the rest remains in steady state. “URE” stands for unhedged interest rate exposure: it inputs the realized return changes after initial revaluations $\{R_{t+s}^b, R_{t+s}^s\}_{s \geq 1}$ leaving all other variables in steady state, including the expected returns that generate the substitution effect.

Figure 13: Decomposition of Aggregate Consumption Response to a Monetary Policy Shock: Asset Returns

out to be lower and higher than expected, respectively. The effect of households’ net nominal position is negligible, and that of the initial capital gains is a persistent 0.1% increment in consumption.²⁴ The remaining component of asset returns is the effect of changes after the initial shock $\{R_{t+s}^b, R_{t+s}^s\}_{s \geq 1}$, net of the intertemporal substitution effect that we analyzed previously. Following Auclert (2019), we denominate this component the URE channel. This channel captures the wealth effect of the announced changes in the sequence of future returns. Since most of our households spend most of their lives accumulating assets (see Figure 4), lower returns work against the affordability of their desired consumption-savings paths. Hence, the URE channel reduces their aggregate initial consumption by 0.3%.

²⁴Part of the reason why the net nominal position channel in our model is so small is that, under the current calibration, the aggregate portfolio share of stocks in households’ assets is more than 0.8, leaving less than 0.2 for bonds. We expect this result to change as we incorporate a richer representation of nominal assets into the model.

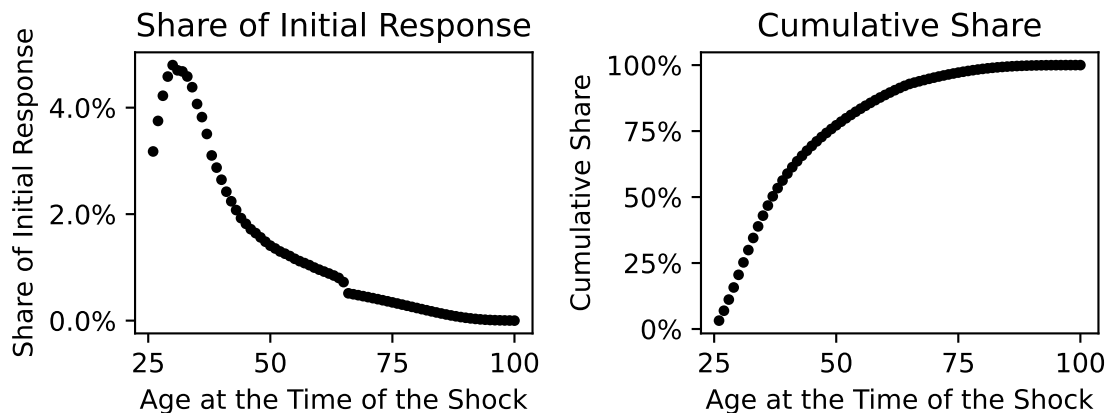
4.2 Monetary Policy Transmission Across the Life Cycle

Having examined the channels of monetary policy transmission in our model and confirmed that the qualitative aggregate conclusions resemble those of previous studies in the HANK literature, we now explore how transmission and its channels vary across agents of different ages. Specifically, this section studies and compares the consumption responses of households that are in different points of their life cycles when the monetary policy shock hits the economy. We compare the responses in both their magnitudes and the mechanisms that generate them. We find that young households are the main contributors to the increase in aggregate consumption and that their response is due in great part to their increased labor income. Intertemporal substitution and realized asset returns increase in importance for the response of older households.

In the first year after the monetary policy shock, more than half of the aggregate consumption response is due to households below the age of 40, and almost its totality is due to households below the age of retirement. Figure 14 depicts the relative contribution of different cohorts to the change in aggregate consumption during the first year of the expansionary monetary shock ($C_t - C_{SS}$). The left panel shows that, after a small initial increase, the contribution of cohorts to the aggregate consumption response steeply declines with their age at the time of the shock: each of the cohorts aged below 30 represents more than 3% of the total response, while none of the retired cohorts represents more than 1%.²⁵ The right panel presents the cumulative shares; it shows that 59% of the initial response in consumption comes from households aged 40 or younger, and that 93% comes from working-age households (aged 65 or younger).

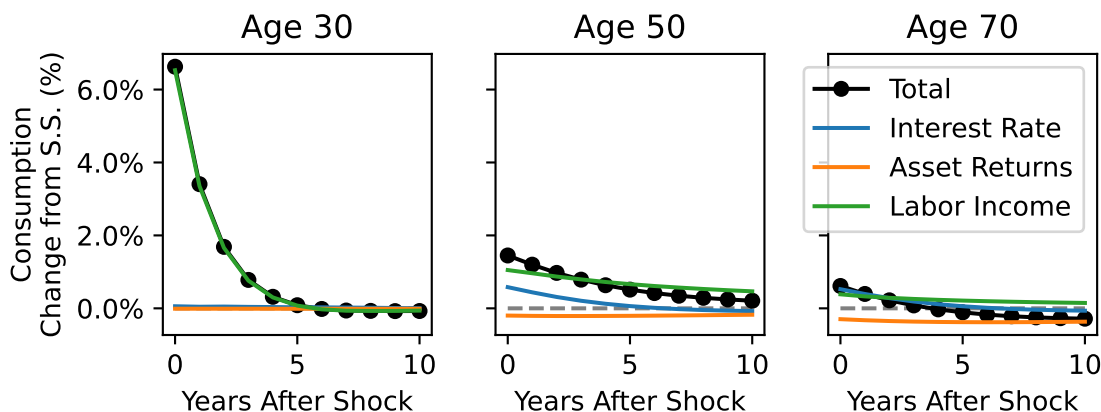
The labor income channel is the main driver of the consumption response for most of the life cycle; for older cohorts, the intertemporal substitution channel gains importance, while the contribution of the asset-returns channel remains small and negative. Figure 15 presents and decomposes the dynamic consumption response of cohorts that were aged 30, 50, and 70 when the monetary policy shock hit. We use the same decomposition that we applied to aggregate consumption: intertemporal substitution, asset returns, and labor income. The 30-year-old households increase their consumption by 6.7% in the first year, driven almost entirely by changes in their labor income—the contribution of intertemporal substitution and asset returns is negligible. Young households in our model behave as “hand-to-mouth” agents that simply consume the increased income that they receive: they have low savings (see Figure 4) and high MPCs (see Figure 8). As households progress through the life cycle and accumulate assets, their consumption response becomes more muted and nuanced. For 50 year olds, consumption increases 1.4% on impact and decays slowly as consumption-smoothing starts to feature: intertemporal substitution represents 0.6 percentage point of this response. Finally, the consumption of 70 year olds increases by only 0.6%, with the labor income channel contributing 0.4 percentage points, intertemporal substitution 0.5, and asset returns -0.3 .

²⁵The size of cohorts is a major factor behind this pattern. Since our model features constant birth and age-specific death rates, the mass of households monotonically declines with age. Figure C.1 depicts the



This figure decomposes the total response in aggregate consumption at the time the monetary policy shock is announced, $C_t - C_{SS}$, into the parts that are due to households of different ages at time t . The left panel displays, for each age, the share of the initial consumption response that comes from the cohort of households of that given age. The right panel presents the cumulative distribution of the shares in the left panel: the share of the response due to households younger than a given age.

Figure 14: Incidence of the Initial Consumption Response Across Cohorts



Each panel presents the response of a different cohort of households to the expansionary monetary policy shock. Cohorts are denominated using the age they were when the shock hit. “Interest Rate” denotes the ex-ante real rate, R^e . “Asset Returns” denotes the realized returns on total assets held by the cohort. “Labor Income” is the total income net of taxes and transfers paid to the cohort. Each line corresponds to the response of the total consumption of the cohort when the deviations in these variables are given to the household sector independently (leaving all other values in steady state) and together (which produces the “Total” response).

Figure 15: Decomposition of Consumption Response to a Monetary Policy Shock by Age

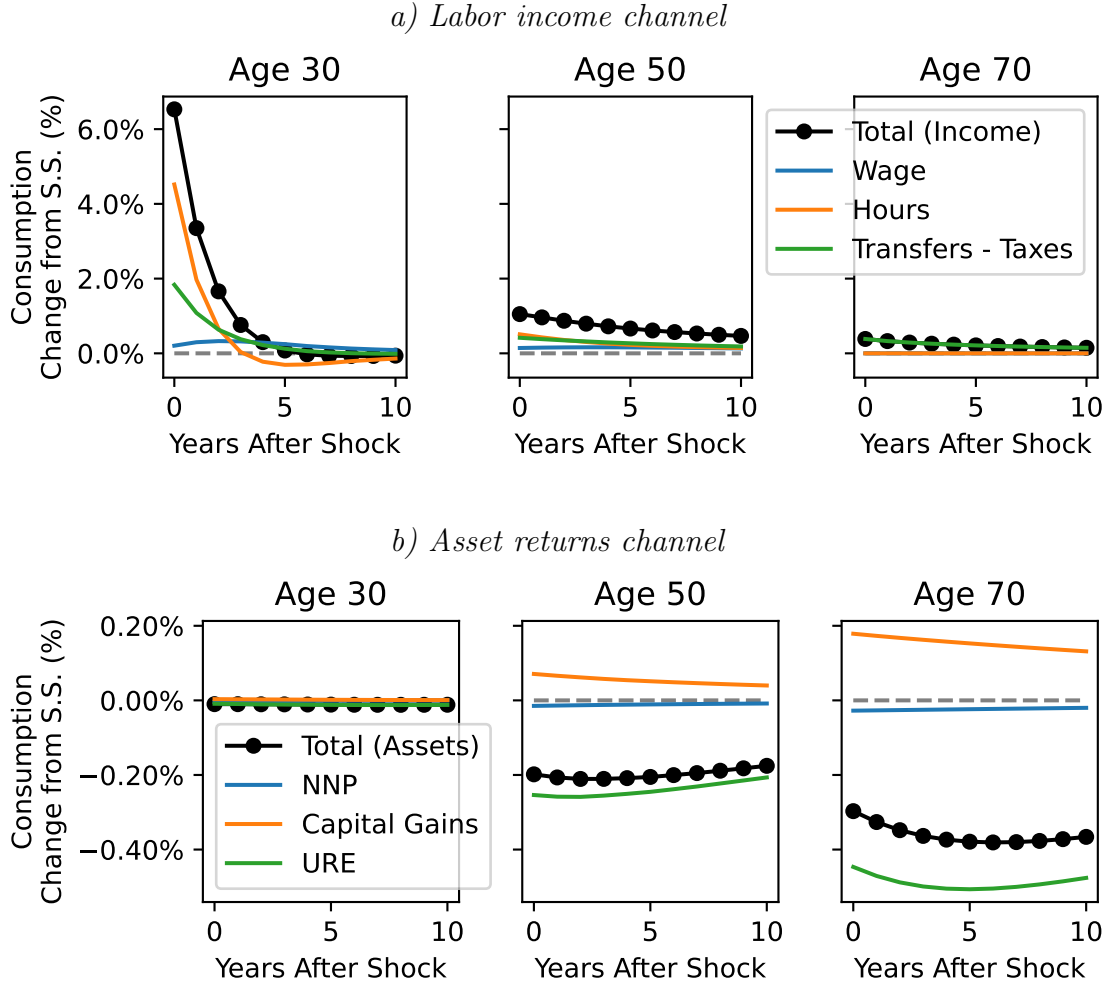
Within the labor-income channel, the fluctuation in hours is the largest contributor to the large consumption response of young households. Panel a) of Figure 16 decomposes the income channel of the consumption response into the contribution of wages, hours, and taxes and transfers. For 30-year-old households, the fluctuation in hours generates 4.5 percentage points of the initial consumption increase, which is more than twice as large as the 1.8 percentage point contribution of fiscal adjustments (transfers minus taxes). For 50 year olds, these two channels become comparable with hours and fiscal adjustments contributing 0.5 percentage point and 0.4 percentage point to initial consumption, respectively. The greater influence of hours for younger agents is due in part to the fact that their hours fluctuate more. The unequal incidence of fluctuations in aggregate demand for hours—which we calibrated to the empirical findings of Jaimovich et al. (2013)—implies that the work hours of 30 year olds increase by 6.4% at the time of the shock, while those of 50 year olds increase only by 4.4%. Wages and hours do not matter for retirees, whose labor income response comes solely from the adjustment of taxes on their retirement benefits.

While monetary policy transmission channels related to asset returns have a small effect on aggregate consumption in our model, the size of this effect increases as households age and, in spite of making them initially wealthier, pushes them to lower their consumption. Panel b) of Figure 16 decomposes the asset returns channel of the response of different cohorts into revaluations in their net nominal position, initial capital gains, and URE. The asset returns channel is inoperative for young hand-to-mouth households. As households age and accumulate more wealth, the effect of the asset returns channel on consumption becomes negative and grows in magnitude; its on-impact effects on consumption are -0.2% for the 50-year-old cohort and -0.3% for the 70-year-old cohort. Consumption contracts even when initial aggregate asset returns turn out to be greater than expected (see Figure 11). This initial wealth revaluation, which combines the net nominal position and capital gains effects, would by itself increase the consumption of the 50-year-old cohort by 0.1% and that of the 70-year-old cohort by 0.2%. However, as pointed out by Auclert (2019), it would be wrong to conclude from this repricing that the new sequence of asset returns leaves these households better off. If, for example, their consumption plan has a higher duration than that of their income, then that plan might become infeasible and they could be left worse off. These dynamic considerations, which indeed turn out to have a negative effect on consumption, are captured by the URE channel. The magnitude of this channel increases with the age of the cohort: its effect on impact is -0.3 percentage point for 50 year olds and -0.4 percentage point for 70 year olds. This is enough to undo the positive wealth revaluation effects.

4.3 The Redistributive Effects of Monetary Policy

We now examine the welfare effects of monetary policy shocks with an emphasis on redistribution. We characterize the households that gain and lose welfare from a monetary easing, and the channels that generate those gains and losses. We use our general equilibrium

demographic distribution of our simulated populations.



Each panel presents the response of a different cohort of households to the expansionary monetary policy shock. Cohorts are denominated using the age they were when the shock hit. “Wage” denotes the real wage, w , “Hours” denotes aggregate labor hours L , “Tax Rate” is $1 - \lambda$ from Equation 6, and “Total (Income)” is the response to these three mechanisms together. “NNP” stands for net nominal position and represents a scenario where only the initial return in bond returns, R_t^b , is passed to the household block. “Capital Gains” isolates the effect of initial equity revaluation. “URE” stands for unhedged interest rate exposure and inputs the realized return changes after initial revaluations $\{R_{t+s}^b, R_{t+s}^e\}_{s \geq 1}$ leaving all other variables in steady state, including the expected returns that generate the substitution effect. “Total (Assets)” aggregates the effects of NNP, Capital Gains, and URE.

Figure 16: The Labor-Income and Asset>Returns Channels by Age

model to account for the many different channels that are relevant to households and use their value functions to find heterogeneous welfare impacts that incorporate considerations such as uncertainty and binding constraints.

Welfare metric. We first introduce the metric that we use to quantify welfare changes. From the perspective of a household, a monetary policy shock at time t is a change in the sequence of macroeconomic aggregates that it observes and expects $\{w_{t+s}, L_{t+s}, R_{t+s}^b, R_{t+s}^s, \lambda_{t+s}\}_{s \geq 0}$. Denoting with V_a^{SS} the value function of the household before the shock is announced—when it expects all aggregates to remain at their steady-state values—and with $V_{a,t}^*$ its value function right after the announcement, the effect of the shock on the welfare of household i is

$$V_{a,t}^*(\alpha_i, z_i, a_{i,t-1}) - V_a^{SS}(\alpha_i, z_i, a_{i,t-1}).$$

To make this quantity interpretable, we use a first-order approximation to find its equivalent variation and re-scale it by the planned consumption of the household:

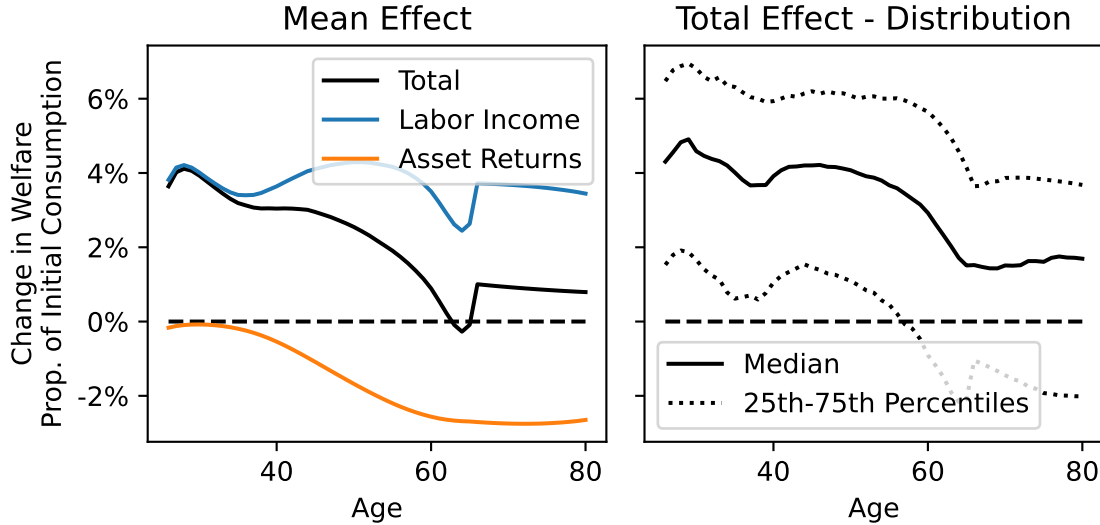
$$\Delta_i \equiv \frac{1}{c_a^{SS}(\alpha_i, z_i, a_{i,t-1})} \times \frac{V_{a,t}^*(\alpha_i, z_i, a_{i,t-1}) - V_a^{SS}(\alpha_i, z_i, a_{i,t-1})}{\partial V_a^{SS}(\alpha_i, z_i, a_{i,t-1}) / \partial a_{i,t-1}}. \quad (28)$$

Our metric Δ_i is the transfer that would have taken household i to the same level of welfare as the monetary policy shock takes it, expressed as a fraction of the consumption it planned for the year of the shock.²⁶ A positive Δ_i means household i gained welfare from the monetary policy shock.

Welfare effect by age. The expansionary shock has welfare effects that are, on average, positive for households of most ages and much greater for younger households. The left panel of Figure 17 shows that the mean total effect is as high as to be comparable with 4% of a year’s consumption for the youngest households, and that it declines steeply with age, reaching a size of 0.8% of consumption for 80-year-old households. This pattern is consistent with the findings of Peterman and Sager (2022), who show that lower real interest rates can have large and positive welfare effects in life-cycle economies, and with those of Bielecki et al. (2022), who also find that monetary easings redistribute welfare towards the young (but find welfare losses for the old).

The positive total welfare effect combines a labor income effect that is positive and relatively flat across the life cycle, with a negative effect from asset returns that is negligible for young households but grows in magnitude with age. As in previous sections, we separate the effects of the monetary policy shock on households into a labor-income channel (that encompasses wages, hours, and transfers net of taxes) and an asset-returns channel. An important change is that, since we now examine welfare, we include what we previously called the intertemporal-substitution channel into asset returns: in our quantification of this channel, households observe the full sequence of returns $\{R_{t+s}^b, R_{t+s}^s\}_{t \geq 0}$ and incorporate their

²⁶The true equivalent variation solves $V_{a,t}^*(\alpha_i, z_i, a_{i,t-1}) = V_a^{SS}(\alpha_i, z_i, a_{i,t-1} + x)$. We rely on the approximation $V_a^{SS}(\alpha_i, z_i, a_{i,t-1} + x) \approx V_a^{SS}(\alpha_i, z_i, a_{i,t-1}) + x \times \partial V_a^{SS}(\alpha_i, z_i, a_{i,t-1}) / \partial a_{i,t-1}$.



Welfare effects are the change in expected discounted utility from period t onwards. We express these changes as an equivalent monetary transfer, rescaled by the consumption of each household at time t (see Equation 28). The left panel presents the mean total effect, and the mean effects due to the labor-income and asset-returns channels for households of every age. The right panel presents the 25th, 50th, and 75th percentiles of the distribution of total welfare effects for households of every age.

Figure 17: Welfare Effects of a Monetary Policy Shock Across the Life Cycle

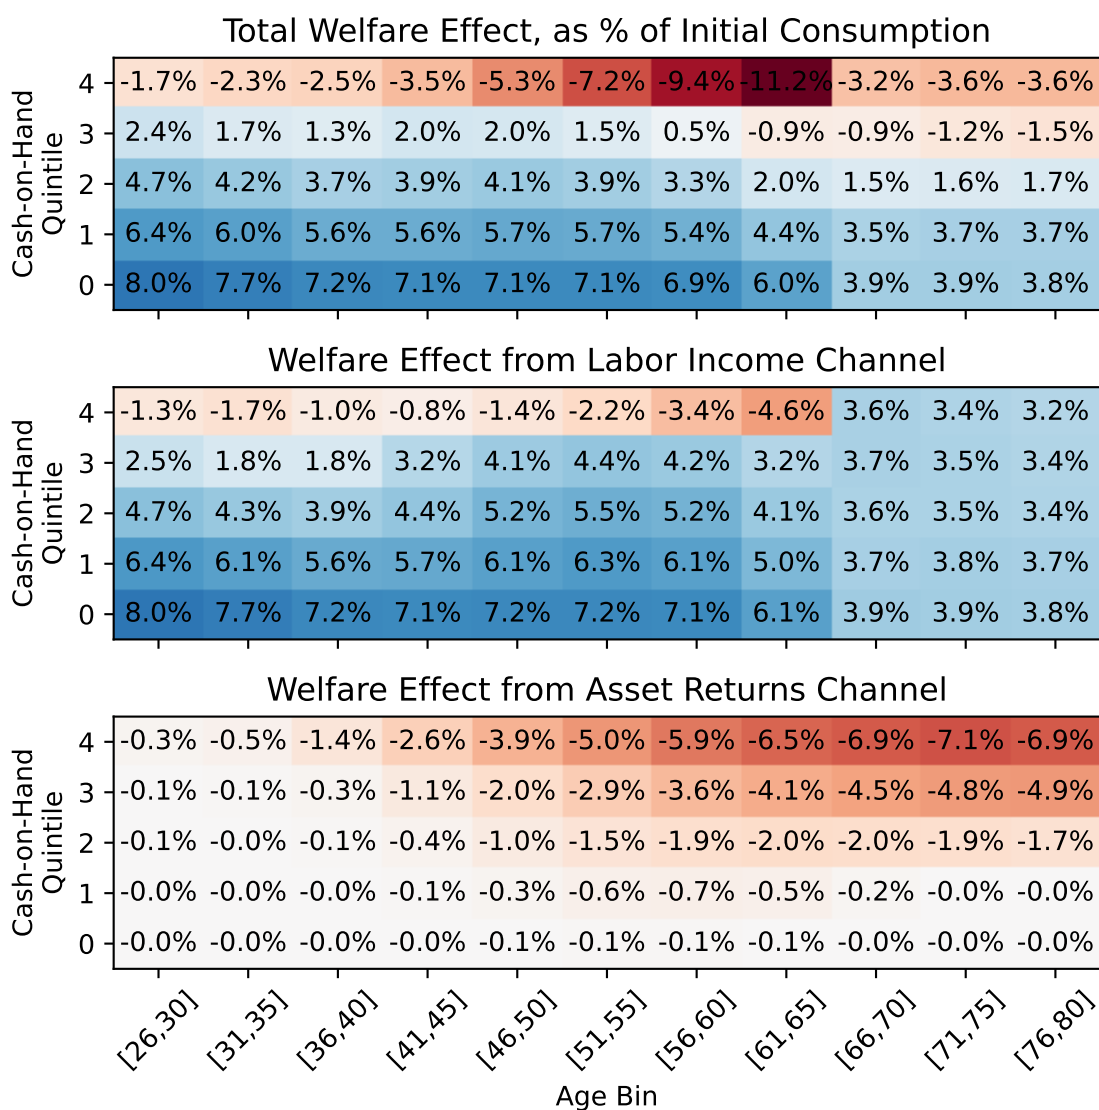
optimal reaction in their expected utility. The other significant consideration is the labor income channel now incorporates the disutility of labor in its welfare effects. The left panel of Figure 17 displays the average welfare effects from the labor-income and asset-returns channel over the life cycle. The labor income of all households increases with the monetary easing because of buoyant labor markets and reduced taxes on non-capital income. This generates a positive and relatively flat average welfare effect of the labor-income channel: it is equivalent to 4.0% of consumption for 30 year olds and 3.5% for 80 year olds. By contrast, the effect of the asset-returns channel is negative and tilted: it is equivalent to -0.1% of consumption for 30 year olds and -2.6% for 80 year olds. Households become exposed to this channel only as they start to accumulate assets or expect that they will in the near future.

Welfare effect by age and wealth. The average welfare effects hide substantial within-age variation that can be larger than cross-age variation. The right panel of Figure 17 depicts the 25th, 50th, and 75th percentiles of the within-age distribution of the total monetary-easing welfare effects for every age. There is a wide range of welfare effects at every age and potential negative effects that were not evident in the average effect. For 30 year olds, the 25th and 75th percentiles of the distribution of welfare effects are, respectively, 1.7% and 6.8% of consumption; for 80 year olds they become -2.0% and 3.7%. In both cases, these ranges are wider than the life-cycle variation of the median welfare effect, which is 4.6% for 30 year olds and 1.7% for 80 year olds. By this measure, while life-cycle differences in the welfare

effects of monetary policy are significant, withing-age heterogeneity is a larger share of their variation. This is notable given that the extant literature on the redistributive effects of monetary policy across generations does not account for within-cohort heterogeneity (Braun & Ikeda, 2021; Bielecki et al., 2022). While age is an important and policy-relevant dimension of heterogeneity, different households of a given age are in vastly different economic shape (see Figures 3 and 4). These differences have a material impact on our conclusions about important quantities such as MPCs and welfare. We now examine heterogeneous welfare effects within age groups.

At any given age, wealthier agents see the greatest welfare losses from the expansionary shock and poorer agents see the greatest gains; this dimension of heterogeneity combines with the age gradient of welfare to produce a wide range of impacts. The top panel of Figure 18 displays the average welfare effect of the expansionary shock for households in different age bins and different quintiles of the bin-specific wealth distribution. For each age bin, the welfare effects monotonically decrease with wealth. The shock redistributes welfare from older and wealthier households to younger and poorer ones. The poorest working-age households perceive gains comparable to 6.0% to 8.0% of their consumption, while the losses of the wealthiest retirees are greater than 3.0% of theirs. Our conclusion about the effects of monetary easing on the distribution of welfare are similar to those of Doepke and Schneider (2006) about the effect of inflation on the distribution of wealth. We now examine how different channels contribute to this welfare redistribution.

The redistributive patterns in the welfare effects from the expansionary shock are the combination of generally positive effects from the labor-income channel, which are greater for young and poor households, and a negative effect from the asset-returns channel that is small for most households but large for the wealthy old. The middle and bottom panels of Figure 18 decompose the average welfare effects of the top panel into its labor-income and asset-returns channels. The middle panel shows that labor-income changes have generally positive welfare impacts. These impacts are the largest for working-age households in the lowest wealth quintile, for whom they range between 6.1% and 8.0% of consumption; they decline steeply as wealth increases. The hours and wage components of the labor-income channel shut down after retirement, resulting in more homogeneous effects that range between 3.2% and 3.9% of consumption and that are due solely to tax reductions. The bottom panel shows that the effects of asset returns are negative at all ages and wealth quintiles. The magnitude of these effects is small for households in the first three wealth quintiles at every age. However, for retirees in the top two quintiles the losses are large: they reach values of up to -7.1% of consumption, which is close in magnitude to the effect of the labor income channel on young and poor households. These losses from the asset returns channel highlight two important points. First, they are the greatest for old and wealthy households, who also have the greatest exposure to the initial revaluation in stocks; this serves to illustrate how initial wealth impacts can differ from welfare impacts, as unhedged exposure to rates of return can undo the effect of initial wealth gains. Second, the welfare losses from this channel are large, even when its contribution to the consumption response was small; this is because the losses are concentrated in wealthy agents that smooth their consumption response over time.



Welfare effects are the change in expected discounted utility from period t onwards. We express these changes as an equivalent monetary transfer, rescaled by the consumption of each household at time t (see Equation 28). We calculate the welfare metric for every household and group them into bins according to their age. We then split them into age bin-specific quintiles of cash-on-hand. For each age bin and wealth quintile, we present the average welfare effect. The top panel considers total welfare effects from the monetary policy shock, while the bottom two isolate the labor-income and asset-returns channels.

Figure 18: Welfare Effects of a Monetary Policy Shock by Age and Wealth

The role of labor supply. Households derive utility from three sources. They value consumption and leaving bequests and dislike working. Following the New Keynesian literature, we assume that hours are determined by labor demand at a common wage per efficiency units. The wage is set by unions that consider the collective welfare of households but individual households have no labor supply choice. This assumption is a simple solution to several well-known problems of HANK models.²⁷ However, it may also compromise the welfare implications of the model, since most households are never on their labor supply curve. See Huo and Ríos-Rull (2020) and Gerke et al. (2023) for detailed discussions of this issue of households “working against their will” in New Keynesian models. We maintain the assumption nevertheless for lack of a better alternative.

As a robustness check, Figure C.4 shows the total welfare effect net of the disutility from labor. As expected, the welfare of working-age households is unambiguously higher, while the welfare of retirees is unaffected. The expansion of labor demand is particularly disliked by rich households close to the retirement age. The general pattern of welfare effects remains the same: young and poor households benefit the most, old and rich households benefit the least, or lose, from monetary easing.

5 Conclusions

This paper has been a first approximation to the study of monetary policy and its transmission channels accounting for the life-cycle transitions of households and for within-cohort heterogeneity. We have demonstrated that a general-equilibrium model that accounts for these facts generates various sharp empirical predictions about, for example, the incidence of the consumption response to monetary shocks across cohorts and the relative importance of different channels for households of different ages. Among these predictions, the fact that most of the consumption response to an expansionary shock comes from the reaction of young households to labor market developments stands out. Another notable prediction is that the revaluation of real assets that occurs after an expansionary shock is not sufficient to compensate their owners for the lower future returns that the shock entails—older wealthy households are harmed by monetary expansions. Our model also delivers predictions about the effect of monetary policy on wealth, income, and consumption inequality that we have omitted from this manuscript but are eager to explore.

We hope that an improved understanding of monetary policy, its transmission channels, and state dependence will come from contrasting the predictions of our model with historical data. This exercise, which we will perform in future versions of this manuscript, will benefit from existing evidence on the distribution of MPCs, the sensitivity of different assets to monetary shocks, and the consumption responses of households of different ages to changes in interest rates, asset prices, and inflation. A richer representation of household balance sheets would be beneficial to these assessments. Housing and mortgages would be a particularly

²⁷First, it improves the co-movement of real wages and asset prices, an important aspect of our exercise (Broer et al., 2020). Second, it shuts down counterfactually large income effects on labor supply without blowing up the multiplier (Auclert et al., 2023).

useful addition to our model.

The exercises that we have performed in this paper also suggest that adding life-cycle considerations to the growing library of HANK models can help address some of their shortcomings. Specifically, we have found that adding a life-cycle motive and a bequest motive for saving appears to help the household block of the model match a broad measure of wealth while retaining a large aggregate MPC, which has been a challenge for existing models. The inclusion of these two motives produces a distribution of savings that is more unequal than those produced by the precautionary motive alone. We believe this fact is behind our achievement. A caveat to this dimension of our analysis is that the current version of our model does not feature purely transitory shocks to income, and that the shocks we include are persistent (not permanent). The precise nature of these shocks has important implications for the precautionary saving motive and MPCs. We aim to evaluate our conclusions under alternative specifications for income risk in future versions of this paper.

References

- Adam, K., & Zhu, J. (2016). Price-Level Changes and the Redistribution of Nominal Wealth across the Euro Area. *Journal of the European Economic Association*, 14(4), 871–906.
- Aguiar, M. A., Bils, M., & Boar, C. (2020, January). *Who Are the Hand-to-Mouth?* 26643.
- Alves, F., Kaplan, G., Moll, B., & Violante, G. L. (2020). A Further Look at the Propagation of Monetary Policy Shocks in HANK. *Journal of Money, Credit and Banking*, 52(S2), 521–559.
- Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., & Tonetti, C. (2020). Long-Term-Care Utility and Late-in-Life Saving. *Journal of Political Economy*, 128(6), 2375–2451.
- Ampudia, M., Georgarakos, D., Slacalek, J., Tristani, O., Vermeulen, P., & Violante, G. (2018, July 18). *Monetary Policy and Household Inequality*.
- Arellano, M., Blundell, R., & Bonhomme, S. (2017). Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. *Econometrica*, 85(3), 693–734.
- Attanasio, O. P., Banks, J., Meghir, C., & Weber, G. (1999). Humps and Bumps in Lifetime Consumption. *Journal of Business & Economic Statistics*, 17(1), 22–35.
- Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6), 2333–2367.
- Auclert, A., Bardóczy, B., & Rognlie, M. (2023). MPCs, MPEs, and Multipliers: A Trilemma for New Keynesian Models. *The Review of Economics and Statistics*, 105(3), 700–712.
- Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica*, 89(5), 2375–2408.
- Auclert, A., Rognlie, M., & Straub, L. (2018, September). *The Intertemporal Keynesian Cross*. 25020.

- Auclert, A., Rognlie, M., & Straub, L. (2020, January). *Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model*. 26647.
- Benhabib, J., & Bisin, A. (2018). Skewed Wealth Distributions: Theory and Empirics. *Journal of Economic Literature*, 56(4), 1261–1291.
- Benhabib, J., Bisin, A., & Luo, M. (2019). Wealth Distribution and Social Mobility in the US: A Quantitative Approach. *American Economic Review*, 109(5), 1623–1647.
- Bielecki, M., Brzoza-Brzezina, M., & Kolasa, M. (2022). Intergenerational Redistributive Effects of Monetary Policy. *Journal of the European Economic Association*, 20(2), 549–580.
- Bilbiie, F. O. (2018, January 1). *Monetary Policy and Heterogeneity: An Analytical Framework*. Retrieved March 14, 2024, from <https://papers.ssrn.com/abstract=3106805>
- Braun, R., & Ikeda, D. (2021, August 20). *Monetary Policy over the Life Cycle* (FRB Atlanta Working Paper No. 2021-20a). Federal Reserve Bank of Atlanta.
- Broer, T., Harbo Hansen, N.-J., Krusell, P., & Öberg, E. (2020). The New Keynesian Transmission Mechanism: A Heterogeneous-Agent Perspective. *The Review of Economic Studies*, 87(1), 77–101.
- Bullard, J. B., DiCecio, R., Singh, A., & Suda, J. (2023, July 19). *Optimal Macroeconomic Policies in a Heterogeneous World*.
- Cagetti, M. (2003). Wealth Accumulation Over the Life Cycle and Precautionary Savings. *Journal of Business & Economic Statistics*, 21(3), 339–353.
- Carroll, C. D. (1992). The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, 1992(2), 61–156.
- Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis*. *The Quarterly Journal of Economics*, 112(1), 1–55.
- Carroll, C. D. (2002, March 30). Why Do the Rich Save So Much? In J. B. Slemrod (Ed.), *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*. Harvard University Press.
- Carroll, C. D., Crawley, E., Slacalek, J., Tokuoka, K., & White, M. N. (2020). Sticky Expectations and Consumption Dynamics. *American Economic Journal: Macroeconomics*, 12(3), 40–76.
- Carroll, C. D., Slacalek, J., Tokuoka, K., & White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3), 977–1020.
- Castañeda, A., Díaz-Giménez, J., & Ríos-Rull, J.-V. (2003). Accounting for the U.S. Earnings and Wealth Inequality. *Journal of Political Economy*, 111(4), 818–857.
- Christiano, L. J., Eichenbaum, M., & Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1), 1–45.
- Clark, K. B., & Summers, L. H. (1981). Demographic Differences in Cyclical Employment Variation. *The Journal of Human Resources*, 16(1), 61–79.
- Crawley, E., & Theloudis, A. (2024, April 18). *Income Shocks and their Transmission into Consumption*. arXiv: 2404.12214 [econ, q-fin].

- De Nardi, M. (2004). Wealth Inequality and Intergenerational Links. *The Review of Economic Studies*, 71(3), 743–768.
- De Nardi, M. (2015, April). *Quantitative Models of Wealth Inequality: A Survey*. 21106.
- De Nardi, M., French, E., & Jones, J. B. (2010). Why Do the Elderly Save? The Role of Medical Expenses. *Journal of Political Economy*, 118(1), 39–75.
- Doepke, M., & Schneider, M. (2006). Inflation and the Redistribution of Nominal Wealth. *Journal of Political Economy*, 114(6), 1069–1097.
- Fagereng, A., Gomez, M., Holm, M., Moll, B., & Natvik, G. (2023). Asset-Price Redistribution.
- Fagereng, A., Holm, M. B., & Natvik, G. J. (2021). MPC Heterogeneity and Household Balance Sheets. *American Economic Journal: Macroeconomics*, 13(4), 1–54.
- Fleck, J., Heathcote, J., Storesletten, K., & Violante, G. L. (2021). *Tax and Transfer Progressivity at the US State Level*.
https://www.jofleck.com/files/state_progressivity.pdf
- Gerke, R., Giesen, S., Lozej, M., & Röttger, J. (2023, November 24). *On Household Labour Supply in Sticky-Wage HANK models*.
- Gomme, P., Rogerson, R., Rupert, P., & Wright, R. (2004). The Business Cycle and the Life Cycle. *NBER Macroeconomics Annual*, 19, 415–461.
- Gourinchas, P.-O., & Parker, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, 70(1), 47–89.
- Greenwald, D., Leombroni, M., Lustig, H. N., & Van Nieuwerburgh, S. (2022, July 12). *Financial and Total Wealth Inequality with Declining Interest Rates*.
- Grigsby, J., Hurst, E., & Yildirmaz, A. (2021). Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data. *American Economic Review*, 111(2), 428–471.
- Guerrieri, V., & Lorenzoni, G. (2017). Credit Crises, Precautionary Savings, and the Liquidity Trap*. *The Quarterly Journal of Economics*, 132(3), 1427–1467.
- Guvenen, F., Schulhofer-Wohl, S., Song, J., & Yogo, M. (2017). Worker Betas: Five Facts about Systematic Earnings Risk. *American Economic Review*, 107(5), 398–403.
- He, Z., Liao, G., & Wang, B. (2022, February 21). *What Gets Measured Gets Managed: Investment and the Cost of Capital* (w29775). National Bureau of Economic Research.
- Heathcote, J., Storesletten, K., & Violante, G. L. (2017). Optimal Tax Progressivity: An Analytical Framework*. *The Quarterly Journal of Economics*, 132(4), 1693–1754.
- Huggett, M. (1996). Wealth distribution in life-cycle economies. *Journal of Monetary Economics*, 38(3), 469–494.
- Huo, Z., & Ríos-Rull, J.-V. (2020). Sticky Wage Models and Labor Supply Constraints. *American Economic Journal: Macroeconomics*, 12(3), 284–318.
- Jaimovich, N., Pruitt, S., & Siu, H. E. (2013). The Demand for Youth: Explaining Age Differences in the Volatility of Hours. *American Economic Review*, 103(7), 3022–3044.

- Janssens, E. F., & McCrary, S. (2023, May 17). *Finite-State Markov-Chain Approximations: A Hidden Markov Approach*.
- Jappelli, T., & Pistaferri, L. (2010). The Consumption Response to Income Changes. *Annual Review of Economics*, 2(1), 479–506.
- Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3), 697–743.
- Kaplan, G., & Violante, G. L. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4), 1199–1239.
- Kaplan, G., & Violante, G. L. (2022). The Marginal Propensity to Consume in Heterogeneous Agent Models. *Annual Review of Economics*, 14.
- Koby, Y., & Wolf, C. K. (2020). *Aggregation in Heterogeneous-Firm Models: Theory and Measurement*.
- McKay, A., Nakamura, E., & Steinsson, J. (2016). The Power of Forward Guidance Revisited. *American Economic Review*, 106(10), 3133–3158.
- McKay, A., & Wolf, C. K. (2023). Monetary Policy and Inequality. *Journal of Economic Perspectives*, 37(1), 121–144.
- Nakamura, E., & Steinsson, J. (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models*. *The Quarterly Journal of Economics*, 123(4), 1415–1464.
- Pallotti, F., Paz-Pardo, G., Slacalek, J., Tristani, O., & Violante, G. L. (2023, November 27). *Who Bears the Costs of Inflation? Euro Area Households and the 2021–2022 Shock* (w31896). National Bureau of Economic Research.
- Peterman, W. B., & Sager, E. (2022). Optimal Public Debt with Life Cycle Motives. *American Economic Journal: Macroeconomics*, 14(4), 404–437.
- Quadrini, V., & Rios-Rull, J.-V. (1997, Spring). Understanding the U.S. distribution of wealth. *Federal Reserve Bank of Minneapolis. Quarterly Review - Federal Reserve Bank of Minneapolis*, 21(2), 22–36.
- Rotemberg, J. J. (1982). Monopolistic Price Adjustment and Aggregate Output. *The Review of Economic Studies*, 49(4), 517–531.
- Sabelhaus, J., & Song, J. (2010). The great moderation in micro labor earnings. *Journal of Monetary Economics*, 57(4), 391–403.
- Stachurski, J., & Toda, A. A. (2019). An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity. *Journal of Economic Theory*, 182, 1–24.

Online Appendix

A Appendix to Section 2

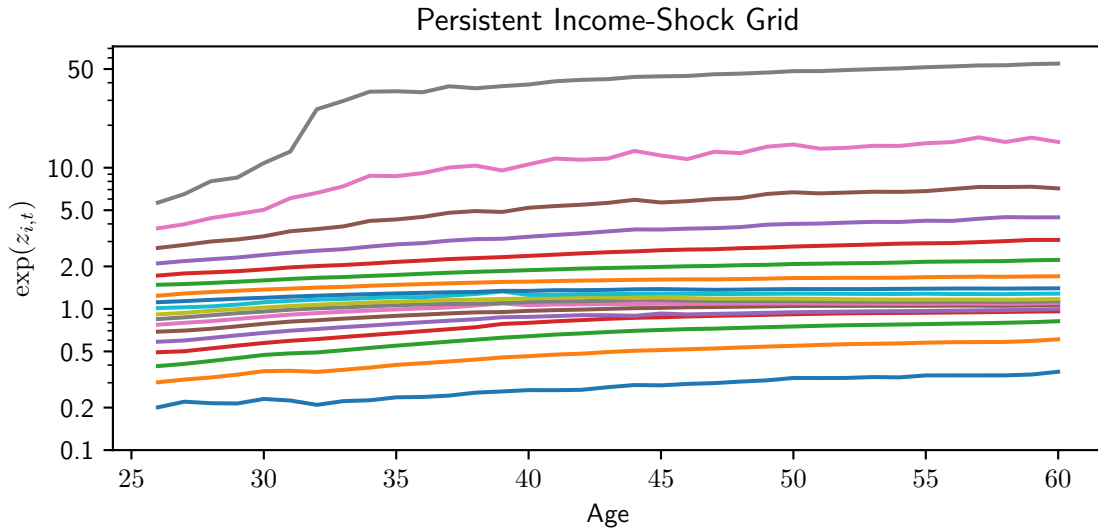
In this appendix, we provide further details on the household block.

A.1 Income Process

The persistent income shock process of Arellano et al. (2017) is estimated using data from the Panel Study of Income Dynamics between 1999 and 2009, and its main innovative features are *nonlinear persistence* and *conditional heterogeneity of higher moments*.

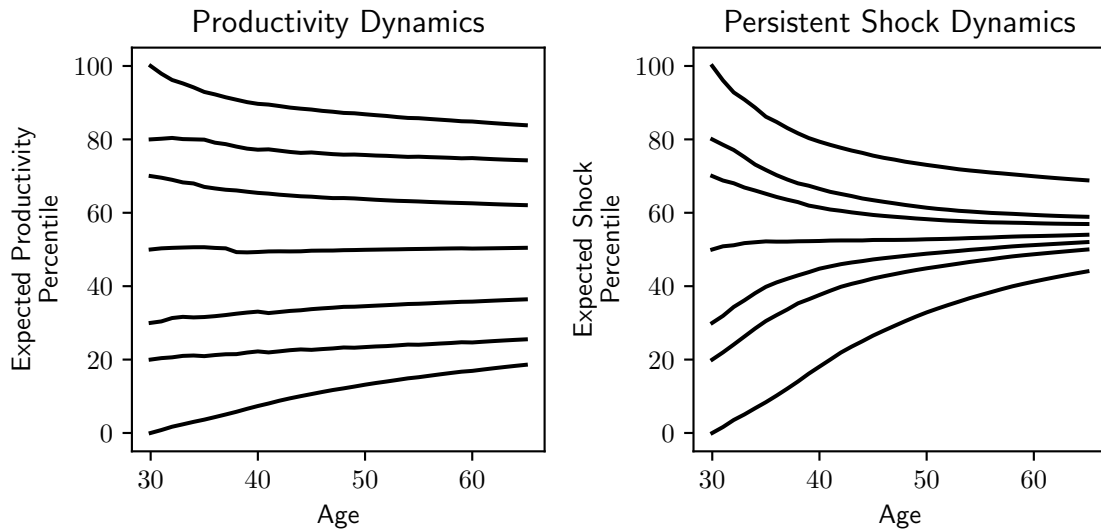
- *Nonlinear persistence* means that the effect of current shocks (z_t) on future shocks (z_{t+1}) can depend on properties of the current shock—for example, its magnitude or sign. This is a departure from canonical unit-root or auto-regressive processes in which shocks of all sizes have the same persistence. The authors find empirical support for this departure in both U.S. and Norwegian data. They develop a measure of persistence that nests the autoregressive coefficient of linear canonical models. This measure of persistence is highest (close to 1.0) when high-earners get positive shocks and low-earners get negative shocks. Negative shocks to high-earners and positive shocks to low-earners are much less persistent (0.6 to 0.8).
- *Conditional heterogeneity of higher moments* means that the distribution of z_{t+1} conditional on z_t is flexible; it is not restricted to a normal distribution with a constant variance as in canonical models. The variance, skewness, etc., of the distribution can vary with z_t . The authors find that, indeed, the skewness of the conditional distribution of z_{t+1} appears to vary with z_t .

In the numerical implementation of our model, we use a discretized version of this income process constructed by Janssens and McCrary (2023). The authors propose an efficient method for obtaining finite-state Markov-chain approximations for continuous stochastic processes like that in Arellano et al. (2017). Their method is efficient in the sense that, for any given number of gridpoints, the approximation minimizes the information loss between the true and the discretized processes. They produce a discretization with 18 age-dependent states and age-dependent transition matrices that very closely reproduces the first four moments of the levels and changes of persistent income shocks in the Arellano et al. (2017) process. We use their discretization, which we display in Figure A.1.



The figure presents, at every age, the set of possible values that our discretized $\exp\{z_{i,t}\}$ can take. These values come from Janssens and McCrary (2023)'s approximation of the income process of Arellano et al. (2017) with 18 points; they were provided to us by the former authors. Note that the figure simply represents the support of the process: it does not contain information about the conditional or unconditional likelihoods of each point.

Figure A.1: Age-Dependent Discretization for Persistent Income Shocks



This figure represents the percentile of the productivity (\tilde{y}) or persistent shock (z) age-specific distribution that a household can expect to be in, given its position in the age-30 distributions. The lines depict the expected trajectories of households that start at the 0, 20th, 30th, 50th, 70th, 80th, and 100th percentiles of the age-30 distributions. Productivity (left panel) combines individual fixed effects α and the persistent shock z . Persistent shock (right panel) is z from the process in Arellano et al. (2017), as discretized by Janssens and McCrary (2023).

Figure A.2: Persistence in Productivity and Shock Dynamics

A.2 Labor Demand

The functional form of the labor demand function in Equation 4 implies that, through a first-order log-approximation

$$\frac{d \ln (\gamma(\mathbf{a}_{i,t}, L_t) / \gamma(\mathbf{a}_{i,t}, L_{ss}))}{d \ln (L_t / L_{ss})} \approx \varepsilon_{\mathbf{a}_{i,t}} + 1 - \frac{\sum_a \tilde{Y}_a \varepsilon_a}{\sum_a \tilde{Y}_a}.$$

Since aggregate fluctuations in labor demand are the only source of age-specific fluctuations in labor demand in the model, this implies that

$$\sigma(\ln(\gamma(\mathbf{a}_{i,t}, L_t))) \approx \sigma(\ln(L_t)) \times \left[\varepsilon_{\mathbf{a}_{i,t}} + 1 - \frac{\sum_a \tilde{Y}_a \varepsilon_a}{\sum_a \tilde{Y}_a} \right]. \quad (29)$$

Equation 29 links the relative volatility of aggregate and age-specific demand for hours. We use this relationship to calibrate our $\{\varepsilon_a\}$ to the cyclical volatilities of hours calculated by Jaimovich et al. (2013). In particular, we take the cyclical volatilities of hours for²⁸

- The 30-39 age group, which is 1.40 percentage points.
- The 60-64 age group, which is 0.73 percentage points.
- The 30-64 age group, which is 1.20 percentage points and which we use as our value for $\sigma(\ln(L_t))$.

We use an affine specification in age for $\{\varepsilon_a\}$,

$$\varepsilon_a = k_0 + k_1 \times a.$$

To estimate (k_0, k_1) , we match the volatility of hours for 35 year olds implied by Equation 29 to the 30-39 estimate of 1.40 and that of 62 year olds to the 60-64 estimate of 0.73. We illustrate the resulting demand function in Figure A.3.


A.3 Calibrating the Equity Share of Assets

This section describes the construction of the equity-share function $\zeta(\mathbf{a}_{i,t}, a_{i,t-1})$ introduced in Section 2.3.

We start by calibrating an intermediate function $\zeta(\tilde{\cdot}, \cdot)$ that captures the relationship of equity shares with age and wealth in the 2019 SCF. We take a flexible functional form $\tilde{\zeta}(\cdot, \cdot; \vartheta)$ with parameters ϑ and estimate ϑ as

$$\hat{\vartheta} = \arg \min_{\vartheta} \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \left\| \tilde{\zeta}(\text{Age}_i, \text{Assets}_i; \vartheta) - \frac{\text{Equity}_i}{\text{Assets}_i} \right\|.$$

²⁸The specific numbers we use come from Table 1, column ‘‘Cyclical volatility.’’



The figure depicts age-specific hours $\{\gamma(a, L)\}_{a=26}^{65}$ for different values of L . Each line represents a different value of L as a multiple of its steady state (80% means aggregate hours are at 80% of their steady state).

Figure A.3: Fluctuations in Age-Specific Working Hours

Our sample \mathcal{I} are respondents in the SCF who report strictly positive income and non-negative assets and who are above the age of 21. For the functional form of $\tilde{\zeta}(\cdot, \cdot; \vartheta)$, we choose a feed-forward neural network with two inputs (age and assets), a single output (the equity share), and a single hidden layer with eight neurons.²⁹ The simple structure of the network balances our goals of capturing the principal non-linearities in the data and preserving a smooth function.

Figure A.4 displays the estimated equity-share function $\tilde{\zeta}(\cdot, \cdot)$. Our estimated function captures various well-known empirical features, like a large fraction of the population not investing in equities and a steep relationship between the equity share and wealth. We construct the function that we use in our model solution by evaluating the estimated net $\tilde{\zeta}(\cdot, \cdot; \hat{\vartheta})$ on our asset grid at every age. Figure A.4 also presents this evaluated function.

In Section 3.2, we target a steady-state aggregate equity share for households, Ξ_{SS} . To satisfy this constraint, we rescale the intermediate $\tilde{\zeta}(\cdot, \cdot)$ function to obtain our actual equity share function $\zeta(\cdot, \cdot)$. We define $\zeta(\cdot, \cdot)$ as

$$\zeta(\mathbf{a}_{i,t}, a_{i,t-1}) \equiv \tilde{\zeta}(\mathbf{a}_{i,t}, a_{i,t-1}) \times \Xi_{SS} \times \frac{\int a_{i,t-1} dD_{ss}}{\int \tilde{\zeta}(\mathbf{a}_{i,t}, a_{i,t-1}) \times a_{i,t-1} dD_{ss}}.$$

²⁹We use ReLU activations for the input and hidden layers, and a sigmoid activation for the output layer to impose the restriction that shares must be in the unit interval.




Figure A.4: Share of Assets in Equity, Estimated Function

B New Keynesian Closure

This section contains the setup and detailed derivations of the production and labor blocks.

B.1 Production Block

The production block has two types of firms: a representative final goods producer and a unit mass of intermediate goods producers. The role of this final goods firm is to provide a microfoundation for the demand curve faced by the monopolists. The role of the intermediate goods firms is to pin down labor demand, investment, and inflation via a New Keynesian Phillips curve.

B.1.1 Final Good Producer

Setup. There is a representative firm that buys a continuum of intermediate goods $\{y_{jt}\}$ and turns them into the homogeneous final good Y_t via a constant elasticity of substitution (CES) production function with elasticity $\epsilon_p > 1$. Let the price of the final good be P_t . The profit maximization problem of the firm is

$$\max_{Y_t, \{y_{jt}\}} P_t Y_t - \int_0^1 p_{jt} y_{jt} dj \quad \text{s.t.} \quad Y_t = \left(\int_0^1 y_{jt}^{\frac{\epsilon_p - 1}{\epsilon_p}} dj \right)^{\frac{\epsilon_p}{\epsilon_p - 1}}$$

Derivation. Substitute the constraint

$$\max_{\{y_{jt}\}} P_t \left(\int_0^1 y_{jt}^{\frac{\epsilon_p-1}{\epsilon_p}} dj \right)^{\frac{\epsilon_p}{\epsilon_p-1}} - \int_0^1 p_{jt} y_{jt} dj$$

The first order condition (FOC) for y_{jt} is

$$0 = P_t \left(\int_0^1 y_{jt}^{\frac{\epsilon_p-1}{\epsilon_p}} dj \right)^{\frac{1}{\epsilon_p-1}} y_{jt}^{\frac{-1}{\epsilon_p}} - p_{jt} = P_t Y_t^{\frac{1}{\epsilon_p}} y_{jt}^{\frac{-1}{\epsilon_p}} - p_{jt}$$

which implies demand curves

$$y_{jt} = \left(\frac{p_{jt}}{P_t} \right)^{-\epsilon_p} Y_t \quad (30)$$

The final good firm is competitive and makes zero profits. This implies

$$P_t Y_t = \int_0^1 p_{jt} y_{jt} dj = \int_0^1 p_{jt} \left(\frac{p_{jt}}{P_t} \right)^{-\epsilon_p} Y_t dj = P_t^{\epsilon_p} Y_t \int_0^1 p_{jt}^{1-\epsilon_p} dj$$

and, hence, the price index is

$$P_t = \left(\int_0^1 p_{jt}^{1-\epsilon_p} dj \right)^{\frac{1}{1-\epsilon_p}} \quad (31)$$

B.1.2 Intermediate goods producers

Setup. There is a unit mass of firms indexed by $j \in [0, 1]$ who engage in monopolistic competition. They have Cobb-Douglas production function

$$F(k_{jt-1}, l_{jt}) = \Theta_t k_{jt-1}^\alpha l_{jt}^{1-\alpha},$$

face a demand curve

$$y_{jt} = \left(\frac{p_{jt}}{P_t} \right)^{-\epsilon_p} Y_t,$$

and set the price of their product subject to a quadratic price adjustment cost

$$\Xi(p_{jt}, p_{jt-1}) = \frac{\chi_p}{2} \left(\frac{p_{jt}}{p_{jt-1}} - 1 \right)^2.$$

Firms buy a homogeneous investment good at relative price p_t^I and use it to augment their capital stock subject to a quadratic capital adjustment cost

$$\Psi \left(\frac{k_{jt}}{k_{jt-1}} \right) = \frac{\psi}{2} \left(\frac{k_{jt}}{k_{jt-1}} - 1 \right)^2.$$

The profit maximization problem of firm j with states $\Omega_{jt} = \{k_{jt-1}, p_{jt-1}\}$ is

$$V_t(\Omega_{jt-1}) = \max_{k_{jt}, i_{jt}, p_{jt}, l_{jt}, y_{jt}} \frac{p_{jt}}{P_t} y_{jt} - w_t l_{jt} - p_t^I i_{jt} - \Psi \left(\frac{k_{jt}}{k_{jt-1}} \right) k_{jt-1} - \Xi(p_{jt}, p_{jt-1}) Y_t + \mathbb{E}_t \left[\frac{V_{t+1}(\Omega_{jt})}{R_t^e} \right],$$

$$\text{s.t. } k_{jt} = (1 - \delta_k) k_{jt-1} + i_{jt},$$

$$y_{jt} = \left(\frac{p_{jt}}{P_t} \right)^{-\epsilon_p} Y_t,$$

$$y_{jt} = F(k_{jt-1}, l_{jt}).$$

Labor demand derivation. Let's substitute the constraints

$$V_t(\Omega_{jt-1}) = \max_{k_{jt}, p_{jt}, l_{jt}} \frac{p_{jt}}{P_t} F(k_{jt-1}, l_{jt}) - w_t l_{jt} - p_t^I [k_{jt} - (1 - \delta_k) k_{jt-1}]$$

$$- \Psi \left(\frac{k_{jt}}{k_{jt-1}} \right) k_{jt-1} - \Xi(p_{jt}, p_{jt-1}) Y_t + \mathbb{E}_t \left[\frac{V_{t+1}(\Omega_{jt})}{R_t^e} \right] + \eta_{jt} \left[F(k_{jt-1}, l_{jt}) - \left(\frac{p_{jt}}{P_t} \right)^{-\epsilon_p} Y_t \right]$$

The FOC for l_{jt} is

$$0 = \frac{p_{jt}}{P_t} \partial_L F(k_{jt-1}, l_{jt}) - w_t + \eta_{jt} \partial_L F(k_{jt-1}, l_{jt}).$$

In symmetric equilibrium, this becomes

$$w_t = \partial_L F(K_{t-1}, L_t) (1 + \eta_t) = \partial_L F(K_{t-1}, L_t) m c_t, \quad (32)$$

where $m c_t = 1 + \eta_t$ denotes the real marginal cost (reciprocal of the markup). The marginal product is higher than the wage because the monopolist hires too little labor.

Phillips curve derivation. Note that the partials of price adjustment cost are

$$\partial_{p_{jt}} \Xi(p_{jt}, p_{jt-1}) = \chi_p \left(\frac{p_{jt}}{p_{jt-1}} - 1 \right) \frac{1}{p_{jt-1}},$$

$$\partial_{p_{jt-1}} \Xi(p_{jt}, p_{jt-1}) = -\chi_p \left(\frac{p_{jt}}{p_{jt-1}} - 1 \right) \frac{p_{jt}}{p_{jt-1}^2}.$$

The FOC for p_{jt} is

$$0 = \frac{1}{P_t} F(k_{jt-1}, l_{jt}) - \partial_{p_{jt}} \Xi(p_{jt}, p_{jt-1}) Y_t + \mathbb{E}_t \left[\frac{\partial_p V_{t+1}(\Omega_{jt})}{R_t^e} \right] + \eta_{jt} \left[\epsilon_p \left(\frac{p_{jt}}{P_t} \right)^{-\epsilon_p - 1} \frac{Y_t}{P_t} \right].$$

Let $\pi_t \equiv P_t/P_{t-1}$ denote gross inflation. In symmetric equilibrium, we get

$$0 = \frac{Y_t}{P_t} (1 + \eta_t \epsilon_p) - \chi_p (\pi_t - 1) \frac{Y_t}{P_{t-1}} + \mathbb{E}_t \left[\frac{\partial_p V_{t+1}(\Omega_t)}{R_t^e} \right],$$

$$0 = Y_t (1 + \eta_t \epsilon_p) - \chi_p \pi_t (\pi_t - 1) Y_t + \mathbb{E}_t \left[\frac{\partial_p V_{t+1}(\Omega_t)}{R_t^e} \right] P_t,$$

$$\chi_p \pi_t (\pi_t - 1) = (1 + \eta_t \epsilon_p) + \mathbb{E}_t \left[\frac{\partial_p V_{t+1}(\Omega_t)}{R_t^e} \right] \frac{P_t}{Y_t}.$$

The envelope condition for P_{jt-1} is, using symmetry in the second line,

$$\begin{aligned}\partial_{p_{jt-1}} V_t &= -\partial_{p_{jt-1}} \Xi(p_{jt}, p_{jt-1}) Y_t = \chi_p \left(\frac{p_{jt}}{p_{jt-1}} - 1 \right) \frac{p_{jt}}{p_{jt-1}^2} Y_t, \\ \partial_{p_{t-1}} V_t &= \chi_p \pi_t (\pi_t - 1) \frac{Y_t}{P_{t-1}}.\end{aligned}$$

Combining the FOCs yields

$$\begin{aligned}\chi_p \pi_t (\pi_t - 1) &= (1 + \eta_t \epsilon_p) + \chi_p \mathbb{E}_t \left[\frac{\pi_{t+1} (\pi_{t+1} - 1) Y_{t+1}}{R_t^e} \frac{P_t}{Y_t} \right], \\ \pi_t (\pi_t - 1) &= \frac{\epsilon_p}{\chi_p} \left(\frac{1}{\epsilon_p} + \eta_t \right) + \mathbb{E}_t \left[\frac{\pi_{t+1} (\pi_{t+1} - 1) Y_{t+1}}{R_t^e} \frac{Y_{t+1}}{Y_t} \right].\end{aligned}$$

Substituting $\eta_t = mc_t - 1$ yields the Phillips curve

$$\pi_t (\pi_t - 1) = \frac{\epsilon_p}{\chi_p} \left(mc_t - \frac{\epsilon_p - 1}{\epsilon_p} \right) + \mathbb{E}_t \left[\frac{\pi_{t+1} (\pi_{t+1} - 1) Y_{t+1}}{R_t^e} \frac{Y_{t+1}}{Y_t} \right].$$

A slight rearrangement lets us parameterize the slope of the linearized NKPC, κ_p , directly:

$$\pi_t (\pi_t - 1) = \underbrace{\frac{\epsilon_p}{\chi_p} \frac{\epsilon_p - 1}{\epsilon_p}}_{\kappa_p} \left(\frac{\epsilon_p}{\epsilon_p - 1} mc_t - 1 \right) + \mathbb{E}_t \left[\frac{\pi_{t+1} (\pi_{t+1} - 1) Y_{t+1}}{R_t^e} \frac{Y_{t+1}}{Y_t} \right]. \quad (33)$$

Investment derivation. The FOC for k_{jt} will be

$$\begin{aligned}-p_t^I - \Psi' \left(\frac{k_{jt}}{k_{jt-1}} \right) + \mathbb{E}_t \left[\frac{\partial_{k_{jt}} V_{t+1}(\Omega_{jt})}{R_t^e} \right] &= 0 \\ p_t^I + \Psi' \left(\frac{k_{jt}}{k_{jt-1}} \right) &= \mathbb{E}_t \left[\frac{\partial_{k_{jt}} V_{t+1}(\Omega_{jt})}{R_t^e} \right]\end{aligned}$$

The right-hand side is Q_t by definition, so we have

$$Q_t = p_t^I + \Psi' \left(\frac{k_{jt}}{k_{jt-1}} \right) = p_t^I + \psi \left(\frac{k_{jt}}{k_{jt-1}} - 1 \right)$$

The envelope condition will be

$$\begin{aligned}\partial_{k_{jt-1}} V_t &= \frac{p_{jt}}{P_t} \partial_k F(\cdot) + p_t^I (1 - \delta_k) - \left[-\Psi' \left(\frac{k_{jt}}{k_{jt-1}} \right) \frac{k_{jt}}{k_{jt-1}} + \Psi \left(\frac{k_{jt}}{k_{jt-1}} \right) \right] + \eta_{jt} \partial_k F(\cdot) \\ \partial_{k_{jt-1}} V_t &= \left[\frac{p_{jt}}{P_t} + \eta_{jt} \right] \partial_k F(\cdot) + p_t^I (1 - \delta_k) + \Psi' \left(\frac{k_{jt}}{k_{jt-1}} \right) \frac{k_{jt}}{k_{jt-1}} - \Psi \left(\frac{k_{jt}}{k_{jt-1}} \right) \\ \partial_{k_{jt-1}} V_t &= mc_t \partial_k F(\cdot) + p_t^I (1 - \delta_k) + \Psi' \left(\frac{k_{jt}}{k_{jt-1}} \right) \frac{k_{jt}}{k_{jt-1}} - \Psi \left(\frac{k_{jt}}{k_{jt-1}} \right)\end{aligned}$$

In symmetric equilibrium, we have that

$$\psi \left(\frac{K_t}{K_{t-1}} - 1 \right) = Q_t - p_t^I \quad (34)$$

$$R_t^e Q_t = mc_{t+1} \partial_k F_{t+1}(\cdot) + p_{t+1}^I (1 - \delta_k) - \Psi \left(\frac{K_{t+1}}{K_t} \right) + \frac{K_{t+1}}{K_t} (Q_{t+1} - p_{t+1}^I) \quad (35)$$

B.1.3 Block representation

The retailers maximize their profit, taking input prices as given. Given the constant-returns-to-scale technology, the level of production is not pinned down by prices alone. So we'll consider aggregate demand as an input to the production block in addition to prices. In sum, given the sequences of inputs $\{Y_t, w_t, r_t, p_t^I\}$ and initial condition K_{-1} , the production block returns eight sequences of outputs $\{L_t, K_t, I_t, Q_t, mc_t, \pi_t, p_t, d_t\}$.

- **Production function.** Gives labor:

$$Y_t = F(K_{t-1}, L_t) = \Theta_t K_{t-1}^\alpha L_t^{1-\alpha} \quad (36)$$

- **Labor demand.** Gives marginal cost:

$$mc_t = \frac{w_t}{F_L(K_{t-1}, L_t)} = \frac{1}{1-\alpha} \frac{w_t L_t}{Y_t} \quad (37)$$

- **Phillips curve.** Gives inflation:

$$\pi_t(\pi_t - 1) = \kappa_p \left(\frac{\epsilon_p}{\epsilon_p - 1} mc_t - 1 \right) + \mathbb{E}_t \left[\frac{\pi_{t+1}(\pi_{t+1} - 1) Y_{t+1}}{R_t^e Y_t} \right] \quad (38)$$

- **Marginal Q.** Gives Q_t :

$$Q_t = \psi \left(\frac{K_t}{K_{t-1}} - 1 \right) + p_t^I \quad (39)$$

- **Investment Euler.** Gives capital:

$$R_t^e Q_t = \mathbb{E}_t \left[\alpha \frac{Y_{t+1}}{K_t} mc_{t+1} - p_{t+1}^I \frac{I_{t+1}}{K_t} - \frac{\psi}{2} \left(\frac{K_{t+1}}{K_t} - 1 \right)^2 + \frac{K_{t+1}}{K_t} Q_{t+1} \right] \quad (40)$$

- **Capital law of motion.** Gives investment:

$$I_t = K_t - (1 - \delta_k) K_{t-1} \quad (41)$$

- **Profit.** Gives dividends:

$$d_t = Y_t - w_t L_t - p_t^I I_t - \Psi \left(\frac{K_t}{K_{t-1}} \right) K_{t-1} - \frac{\chi_p}{2} \pi_t^2 Y_t \quad (42)$$

- **Equity price.** Gives equity price:

$$p_t = \mathbb{E}_t \left[\frac{d_{t+1} + p_{t+1}^I}{R_t^e} \right] \quad (43)$$

B.1.4 Calibration strategy.

The calibrated household block pins down $\{R, w, L\}$. We have to ensure that the firm block is consistent with these values. Take the capital-output ratio K/Y and the labor share wL/Y as additional calibration targets.

We don't model the production of investment goods separately. The implicit assumption is that investment goods are produced one-to-one from the final good. Therefore, $p_t^I \equiv 1$. Let steady-state inflation be normalized to $\pi = 0$. Then, equations (39) and (38) imply immediately that

$$Q = 1, \quad mc = \frac{\epsilon_p - 1}{\epsilon_p}. \quad (44)$$

The $\{w, N\}$ inherited from the household block plus the targets for labor share and the capital-output ratio pin down output and capital

$$Y = \left(\frac{wL}{Y}\right)^{-1} wL, \quad K = \left(\frac{K}{Y}\right) Y \quad (45)$$

Next, we can solve for the technology parameters that justify the targeted quantities

$$\alpha = 1 - \frac{wL}{Y} \frac{1}{mc}, \quad \delta_k = 1 + \alpha \frac{Y}{K} mc - R, \quad \Theta = \frac{Y}{K^\alpha L^{1-\alpha}}. \quad (46)$$

Finally, compute investment and flow profits (dividends)

$$I = \delta K, \quad d = Y - wL - I. \quad (47)$$

B.1.5 Block Jacobians

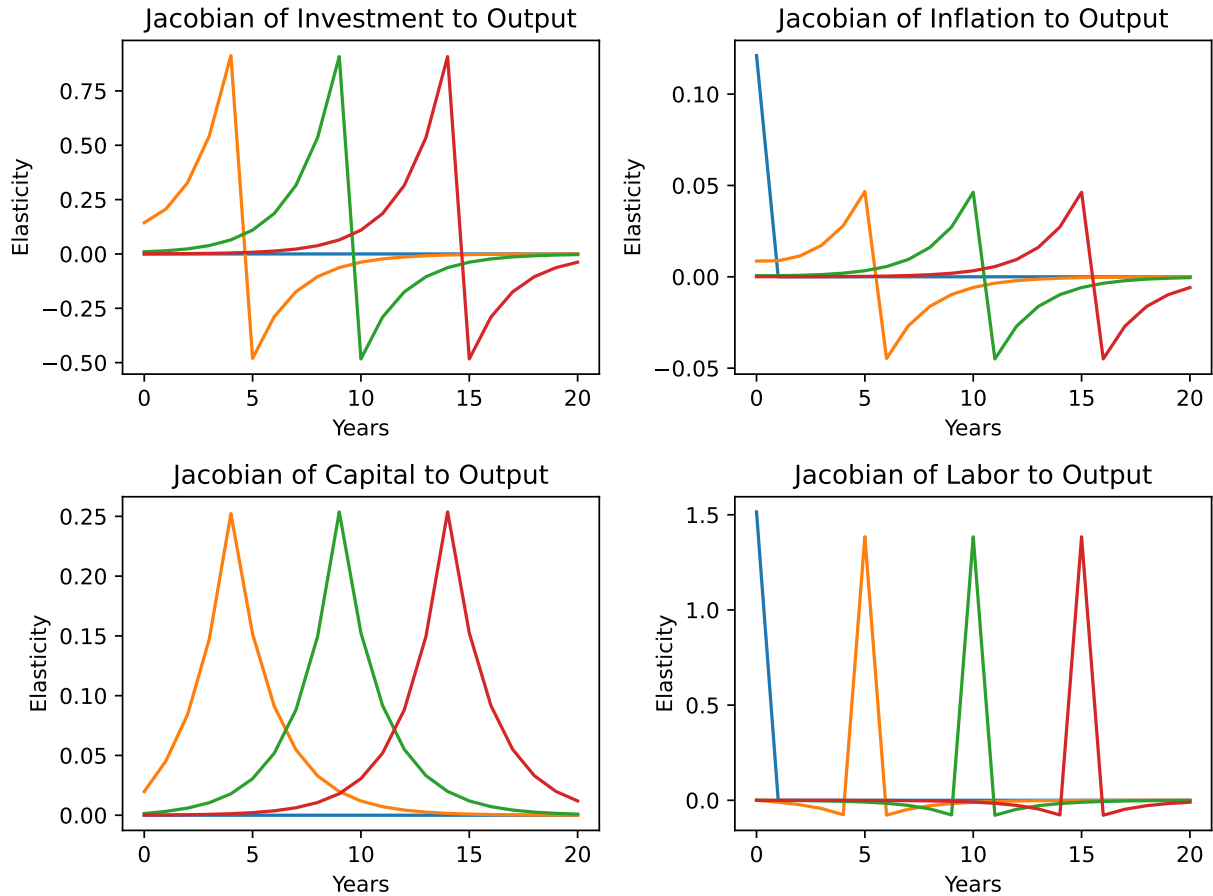


Figure B.1: Jacobians of the Firm Block with Respect to Output

Note: This figure shows $\partial \log O_t / \partial \log Y_s$ for outputs $O \in \{I, K, N, \pi\}$, $t = 0, \dots, 20$, and $s \in \{0, 5, 10, 15\}$. Consider the green line, which shows the responses when firms have to raise output by 1% in period 10. In the absence of capital adjustment costs, firms would change their capital and labor only in period 10. In the presence of convex capital adjustment costs, firms prefer to raise capital gradually in anticipation of the shock. After the shock, they decumulate excess capital gradually. Keeping output constant in periods $t \neq 10$ requires reducing labor demand to offset higher capital. With factor prices being fixed, real marginal cost (not shown) follows a similar path to labor: it's well above steady state in $t = 10$, and weakly below steady state in other periods. A forward-looking Phillips curve implies that inflation rises in anticipation of high marginal cost in period 10 and falls thereafter.




Figure B.2: Jacobians of the Firm Block with Respect to Real Interest Rate

Note: This figure shows $\partial \log O_t / \partial \log R_s^c$ for outputs $O \in \{I, K, N, \pi\}$, $t = 0, \dots, 20$, and $s \in \{0, 5, 10, 15\}$. Consider the green line, which shows the responses when the real interest rate rises by 1% in period 10. This means that firms discount future profits more in periods $t < 10$. Investment and capital falls gradually, given convex adjustment costs. Labor demand rises to offset the effect of lower capital in production. Higher labor costs imply higher real marginal costs in all periods $t \neq 10$. This leads to higher inflation, which is front-loaded due to the strongly forward-looking Phillips curve.

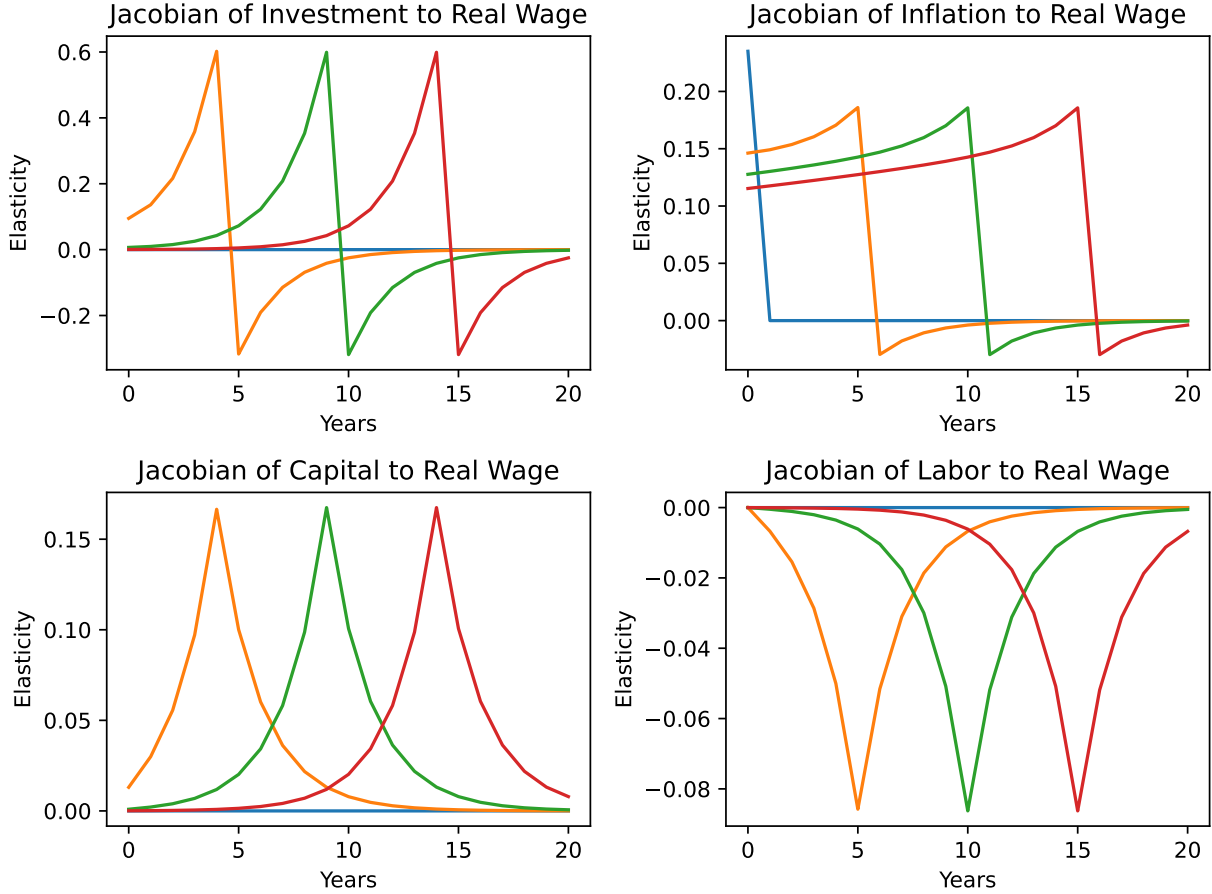


Figure B.3: Jacobians of the Firm Block with Respect to Real Wage

Note: This figure shows $\partial \log O_t / \partial \log w_s$ for outputs $O \in \{I, K, N, \pi\}$, $t = 0, \dots, 20$, and $s \in \{0, 5, 10, 15\}$. Consider the green line, which shows the responses when the real wage rises by 1% in period 10. The optimal capital-labor ratio becomes higher in $t = 10$ only. Due to convex capital adjustment costs, firms accumulate more capital gradually, lowering labor demand in tandem to keep production constant. Real marginal cost rises substantially in period 10 and is weakly below steady state in all other periods. The forward-looking Phillips curve implies a rise in inflation for all $t \leq 10$ and moderate disinflation for $t > 10$.

B.2 Labor Block

The labor block has two types of agents: a representative labor packer and a unit mass of unions. The role of the labor packer is to provide a microfoundation for the demand curves faced by the unions. The role of the unions is to pin down wage inflation via a New Keynesian wage Phillips curve.

B.2.1 Labor packer

Setup. There is a representative firm that buys a continuum of labor services $\{L_{kt}\}$ and turns them into aggregate labor services L_t via a CES production function with elasticity $\epsilon_w > 1$. Let the aggregate nominal wage be W_t , and task-specific nominal wages be $\{W_{kt}\}$.

The profit maximization problem of the labor packer is

$$\max_{L_t, \{L_{kt}\}} W_t L_t - \int_0^1 W_{kt} L_{kt} dk, \quad \text{s.t. } L_t = \left(\int_0^1 L_{kt}^{\frac{\epsilon_w - 1}{\epsilon_w}} dk \right)^{\frac{\epsilon_w}{\epsilon_w - 1}}.$$

Derivation. Substitute the constraint

$$\max_{\{L_{kt}\}} W_t \left(\int_0^1 L_{kt}^{\frac{\epsilon_w - 1}{\epsilon_w}} dk \right)^{\frac{\epsilon_w}{\epsilon_w - 1}} - \int_0^1 W_{kt} L_{kt} dk.$$

The FOC for N_{kt} is

$$0 = W_t \left(\int_0^1 L_{kt}^{\frac{\epsilon_w - 1}{\epsilon_w}} dk \right)^{\frac{1}{\epsilon_w - 1}} L_{kt}^{-\frac{1}{\epsilon_w}} - W_{kt} = W_t L_t^{\frac{1}{\epsilon_w}} L_{kt}^{-\frac{1}{\epsilon_w}} - W_{kt},$$

which implies demand curves

$$L_{kt} = \left(\frac{W_{kt}}{W_t} \right)^{-\epsilon_w} L_t. \quad (48)$$

The labor packer is competitive and makes zero profits. This implies

$$W_t L_t = \int_0^1 W_{kt} L_{kt} dk = \int_0^1 W_{kt} \left(\frac{W_{kt}}{W_t} \right)^{-\epsilon_w} L_t dk = W_t^{\epsilon_w} L_t \int_0^1 W_{kt}^{1 - \epsilon_w} dk,$$

and, hence, the wage index is

$$W_t = \left(\int_0^1 W_{kt}^{1 - \epsilon_w} dk \right)^{\frac{1}{1 - \epsilon_w}}. \quad (49)$$

B.2.2 Unions

Setup. There is a union for every labor service $k \in [0, 1]$ that sets the nominal wage W_{kt} . To ensure that a symmetric equilibrium exists, we assume that every union represents a representative sample of the working-age population. The objective of the unions is to maximize the welfare of working-age households, taking their consumption-saving decisions and the age-specific labor demand schedule as given. There is a quadratic utility cost of adjusting the nominal wage.

The Bellman equation is

$$V_t(W_{k,t-1}) = \int u(c_{i,t}) - v(n_{i,t}) dD_t(\mathbf{a}_{i,t} \leq 65) - \frac{\chi_w}{2} \left(\frac{W_{k,t}}{W_{k,t-1}} - 1 \right)^2 + \beta \mathbb{E}_t [V_{t+1}(W_{k,t})],$$

$$\text{s.t. } L_{k,t} = \left(\frac{W_{k,t}}{W_t} \right)^{-\epsilon} L_t.$$

Derivation. Let's start by rewriting the problem in terms of real wages $w_{k,t} = W_{k,t}/P_t$ and $w_t = W_t/P_t$.

$$V_t(w_{k,t-1}) = \int u(c_{i,t}) - v(n_{i,t}) dD_t(\mathbf{a}_{i,t} \leq 65) - \frac{\chi_w}{2} \left(\pi_t \frac{w_{k,t}}{w_{k,t-1}} - 1 \right)^2 + \beta \mathbb{E}_t [V_{t+1}(w_{k,t})],$$

$$\text{s.t. } L_{k,t} = \left(\frac{w_{k,t}}{w_t} \right)^{-\epsilon} L_t.$$

From now on let's write $d\hat{D}_t = dD_t(\mathbf{a}_{i,t} \leq 65)$ for short. The FOC is

$$0 = \int \left[u'(c_{i,t}) \frac{\partial c_{i,t}}{\partial w_{k,t}} - v'(n_{i,t}) \frac{\partial n_{i,t}}{\partial w_{k,t}} \right] d\hat{D}_t - \chi_w \left(\pi_t \frac{w_{k,t}}{w_{k,t-1}} - 1 \right) \frac{\pi_t}{w_{k,t-1}} + \beta \mathbb{E}_t [V'_{t+1}(w_{k,t})].$$

The envelope condition is

$$V'_t(w_{k,t-1}) = \chi_w \left(\pi_t \frac{w_{k,t}}{w_{k,t-1}} - 1 \right) \pi_t \frac{w_{k,t}}{w_{k,t-1}^2}.$$

Combining these two yields

$$0 = \int \dots d\hat{D}_t - \chi_w \left(\pi_t \frac{w_{k,t}}{w_{k,t-1}} - 1 \right) \frac{\pi_t}{w_{k,t-1}} + \beta \chi_w \left(\pi_{t+1} \frac{w_{k,t+1}}{w_{k,t}} - 1 \right) \pi_{t+1} \frac{w_{k,t+1}}{w_{k,t}^2}.$$

In symmetric equilibrium, all unions set the same wage. Let's define wage inflation $\pi_t^w \equiv \pi_t w_t / w_{t-1}$. Then, we can write

$$(\pi_t^w - 1) \pi_t^w = \frac{w_t}{\chi_w} \int \dots d\hat{D}_t + \beta (\pi_{t+1}^w - 1) \pi_{t+1}^w. \quad (50)$$

Unpacking the integral. Let's start with the disutility of labor. For worker i represented by union k , labor demand is

$$n_{i,t} = \gamma(\mathbf{a}_{i,t}, L_{k,t}) = \gamma \left(\mathbf{a}_{i,t}, \left(\frac{w_{k,t}}{w_t} \right)^{-\epsilon_w} L_t \right).$$

Thus, the partial is

$$\frac{\partial n_{i,t}}{\partial w_{k,t}} = \frac{\partial \gamma(\mathbf{a}_{i,t}, L_{k,t})}{\partial L_{k,t}} \frac{\partial L_{k,t}}{\partial w_{k,t}} = \frac{\partial \gamma(\mathbf{a}_{i,t}, L_{k,t})}{\partial L_{k,t}} \left(-\epsilon_w \frac{L_{k,t}}{w_{k,t}} \right), \quad (51)$$

where, given the functional form in Equation 4,

$$\begin{aligned} \frac{\partial \gamma(\mathbf{a}_{i,t}, L_{k,t})}{\partial L_{k,t}} &= \frac{\partial}{\partial L_{k,t}} \left[L_{k,t} \times \frac{\left(\frac{L_{k,t}}{L_{ss}} \right)^{\epsilon_{a_{i,t}}}}{\sum_{\mathbf{a}} \tilde{Y}_{\mathbf{a}} \times \left(\frac{L_{k,t}}{L_{ss}} \right)^{\epsilon_{\mathbf{a}}}} \right] \\ &= \frac{\gamma(\mathbf{a}_{i,t}, L_{k,t})}{L_{k,t}} - \frac{\sum_{\mathbf{a}} \tilde{Y}_{\mathbf{a}} \times (\epsilon_{\mathbf{a}} - \epsilon_{a_{i,t}}) \times \left(\frac{L_{k,t}}{L_{ss}} \right)^{\epsilon_{\mathbf{a}} - \epsilon_{a_{i,t}}}}{\left[\sum_{\mathbf{a}} \tilde{Y}_{\mathbf{a}} \times \left(\frac{L_{k,t}}{L_{ss}} \right)^{\epsilon_{\mathbf{a}} - \epsilon_{a_{i,t}}} \right]^2}. \end{aligned}$$

Note that the envelope theorem applied to the household problem implies that we can evaluate indirect utility as if marginal changes in income are consumed fully. So, instead of consumption, we can take the partial of real, post-tax income:

$$z_{i,t}(w_{k,t}) \equiv T(\tilde{y}_{i,t} \times w_{k,t} \times n_{i,t}), \quad (52)$$

which is

$$\frac{\partial z_{i,t}}{\partial w_{k,t}} = T'(y_{i,t})\tilde{y}_{i,t} \left[n_{i,t} + w_{k,t} \frac{\partial n_{i,t}}{\partial w_{k,t}} \right] \quad (53)$$

$$\frac{\partial z_{i,t}}{\partial w_{k,t}/w_t} = T'(y_{i,t})y_{i,t} \left[1 + \frac{\partial n_{i,t}/n_{i,t}}{\partial w_{k,t}/w_t} \right] = T'(y_{i,t})y_{i,t} \left[1 - \epsilon_w \frac{L_t}{n_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} \right] \quad (54)$$

Plug these back into (50) to get

$$\begin{aligned} (\pi_t^w - 1) \pi_t^w &= \frac{1}{\chi_w} \left[\int L_t v'(n_{i,t}) \epsilon_w \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} d\hat{D}_t \right. \\ &\quad \left. - \int u'(c_{it}) T'(y_{i,t}) y_{i,t} \left(\epsilon_w \frac{L_t}{n_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} - 1 \right) d\hat{D}_t \right] \\ &\quad + \beta (\pi_{t+1}^w - 1) \pi_{t+1}^w. \end{aligned} \quad (55)$$

Scaling the Phillips curve. Let's define distributional aggregates

$$u'(C_t^*) = \int u'(c_{it}) T'(y_{i,t}) y_{i,t} \left(\epsilon_w \frac{N_t}{n_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, \tilde{N}_t)}{\partial N_t} - 1 \right) d\hat{D}_t, \quad (56)$$

$$v'(N_t^*) = \int L_t v'(n_{i,t}) \epsilon_w \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} d\hat{D}_t. \quad (57)$$

Then, the nonlinear Phillips curve becomes

$$(\pi_t^w - 1) \pi_t^w = \frac{1}{\chi_w} \left[v'(L_t^*) - u'(C_t^*) \right] + \beta (\pi_{t+1}^w - 1) \pi_{t+1}^w. \quad (58)$$

The New Keynesian literature typically works with loglinearized Phillips curves, and calibrate the slope based on the frequency of price or wage adjustments. In order to follow that strategy, it's useful to loglinearize (58). Let hatted variables denote log-deviations from steady state, assuming that gross inflation in steady state is 1, we get

$$\hat{\pi}_t^w = \frac{1}{\chi_w} \left[v''(N^*) N^* \hat{N}_t^* - u''(C^*) C^* \hat{C}_t^* \right] + \beta \hat{\pi}_{t+1}^w$$

Using that $v'(N^*) = u'(C^*)$, we can write this as

$$\hat{\pi}_t^w = \frac{v'(N^*)}{\chi_w} \left[\frac{v''(N^*) N^*}{v'(N^*)} \hat{N}_t^* - \frac{u''(C^*) C^*}{u'(C^*)} \hat{C}_t^* \right] + \beta \hat{\pi}_{t+1}^w, \quad (59)$$

$$\hat{\pi}_t^w = \underbrace{\frac{v'(N^*)}{\chi_w}}_{\kappa_w} \left[\frac{1}{\nu} \hat{N}_t^* + \rho \hat{C}_t^* \right] + \beta \hat{\pi}_{t+1}^w, \quad (60)$$

where $\nu > 0$ is the Frisch elasticity, and $\rho > 0$ is relative risk aversion. Equation (60) has the form of a textbook New Keynesian wage Phillips curve with slope $\kappa_w > 0$. In a Calvo model where unions can adjust wages with probability ξ_w , the slope is

$$\kappa_w = \frac{1}{1 + \Gamma_w} \frac{[1 - \beta(1 - \xi_w)]\xi_w}{1 - \xi_w}, \quad (61)$$

where $\Gamma_w \geq 0$ captures real rigidities. This formula allows us to calibrate the slope of the Phillips curve based on frequency of wage adjustment in micro data.

As a final step, note that we can rewrite the nonlinear Phillips curve (58) as

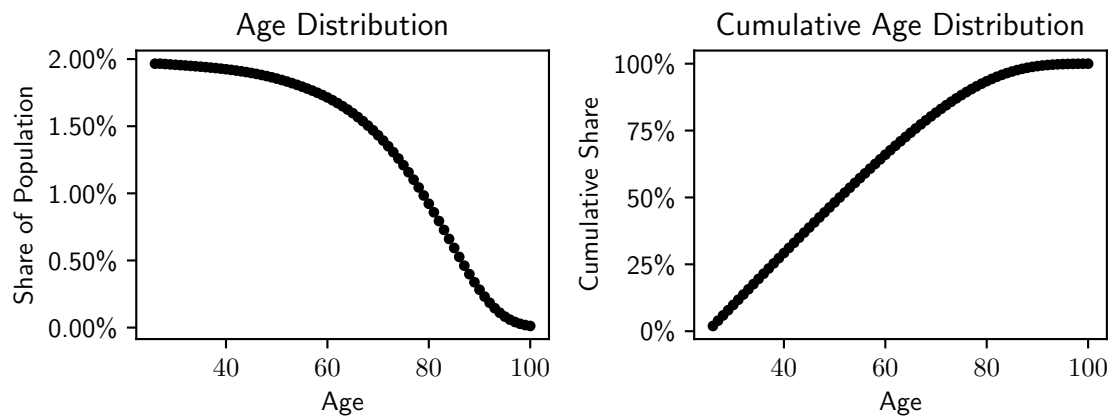
$$(\pi_t^w - 1) \pi_t^w = \kappa_w \left(\frac{v'(N_t^*)}{u'(C_t^*)} - 1 \right) + \beta (\pi_{t+1}^w - 1) \pi_{t+1}^w \quad (62)$$

which is equivalent up to first-order, but is conveniently parameterized with κ_w .

Finally, we note that

$$\begin{aligned} \lim_{\epsilon_w \rightarrow \infty} \frac{v'(N_t^*)}{u'(C_t^*)} &= \lim_{\epsilon_w \rightarrow \infty} \frac{\int L_t v'(n_{i,t}) \epsilon_w \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} d\hat{D}_t}{\int u'(c_{it}) T'(y_{i,t}) y_{i,t} \left(\epsilon_w \frac{N_t}{n_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, \tilde{N}_t)}{\partial N_t} - 1 \right) d\hat{D}_t} \\ &= \lim_{\epsilon_w \rightarrow \infty} \frac{\int L_t v'(n_{i,t}) \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} d\hat{D}_t}{\int u'(c_{it}) T'(y_{i,t}) y_{i,t} \left(\frac{N_t}{n_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, \tilde{N}_t)}{\partial N_t} - \frac{1}{\epsilon_w} \right) d\hat{D}_t} \\ &= \frac{\int L_t v'(n_{i,t}) \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} d\hat{D}_t}{\int u'(c_{it}) T'(y_{i,t}) y_{i,t} \frac{N_t}{n_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, \tilde{N}_t)}{\partial N_t} d\hat{D}_t}. \end{aligned}$$


C Additional Tables and Figures




This figure depicts the age distribution of our simulated populations. These come from sequentially applying survival probabilities in the SSA life tables to an initial mass of agents at age 26. The left panel displays, for each age, the share of the population that has that age at any time. The right panel presents the cumulative distribution of the shares in the left panel: the share of households younger than a given age.

Figure C.1: Age Distribution of Households

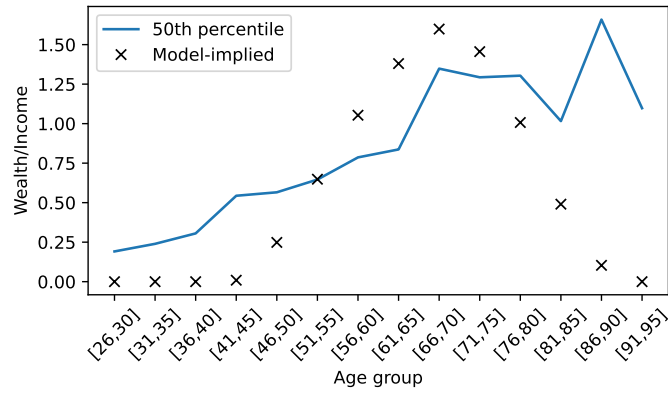
a) Fit of the “No Inequality” model



b) Fit of the “No Bequests” model

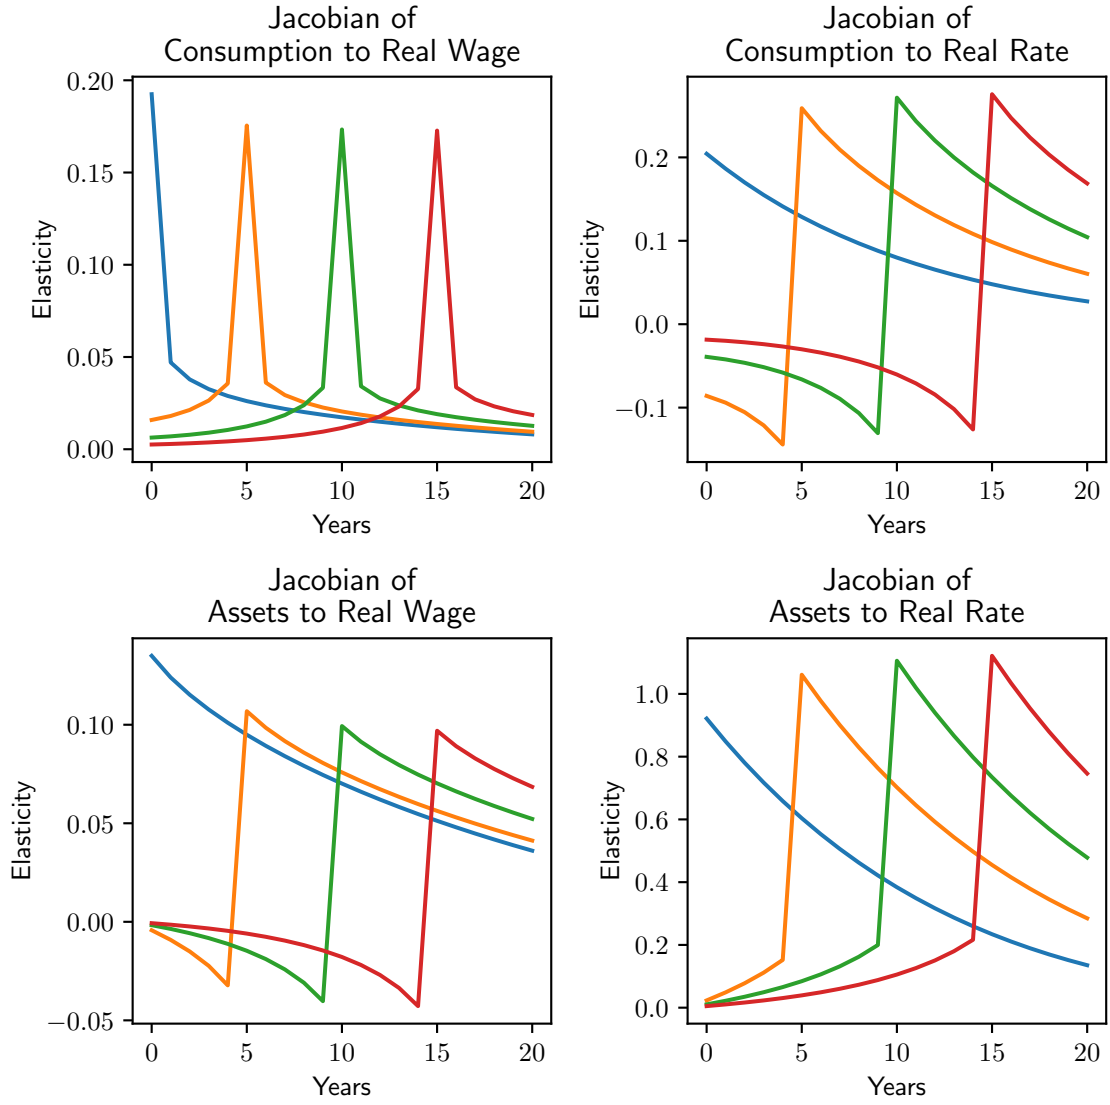


c) Fit of the “No Inequality or Bequests” model



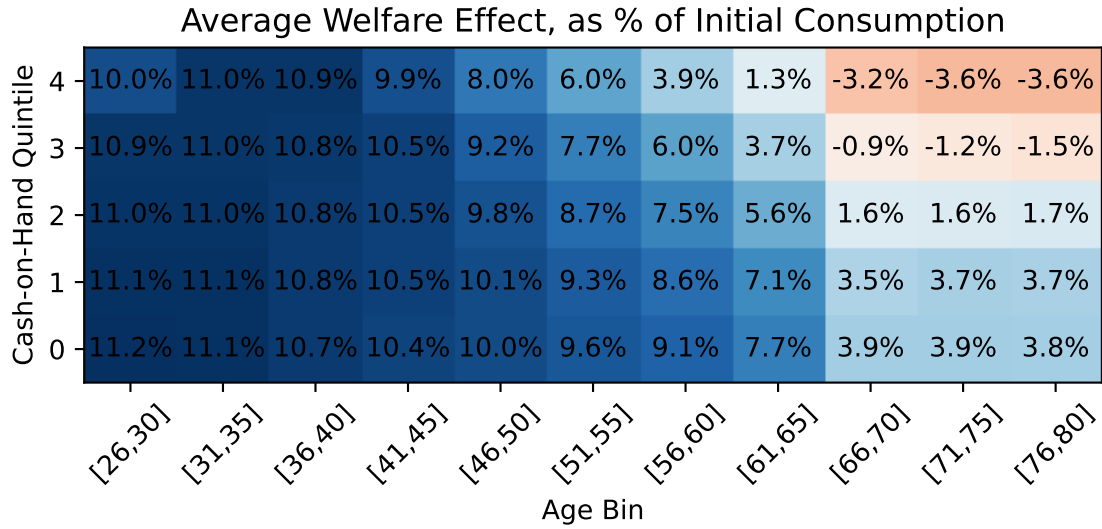
See Section 2 for the definitions of each model specification.

Figure C.2: Fit of Wealth Ratios in Alternative Models



This figure shows $\partial \log O_t / \partial \log I_s$ for outputs $O \in \{C, A\}$, $O \in \{w, R\}$, $t = 0, \dots, 20$, and $s \in \{0, 5, 10, 15\}$. “Real Rate” jacobians consider simultaneous and equivalent increases to expected rates R_{t+s-1}^e and realized rates $\{R_{t+s}^b, R_{t+s}^s\}$.

Figure C.3: Jacobians of the Household Block



Welfare effects are the change in expected discounted utility from period t onwards. We express these changes as an equivalent monetary transfer, rescaled by the consumption of each household at time t (see Equation 28). We calculate the welfare metric for every household and group them into bins according to their age. We then split them into age bin-specific quintiles of cash-on-hand. For each age bin and wealth quintile, we present the average welfare effect.

Figure C.4: Welfare Effects of a Monetary Policy Shock Net of Disutility of Labor Supply