

HANK Comes of Age

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

¹Federal Reserve Board

First version: July, 2024
This version: November, 2024

Abstract

We study monetary policy in a heterogeneous agent New Keynesian model that represents the life cycle of households. The model matches the distribution of labor income and wealth by age. It also produces a realistic distribution of MPCs. Monetary policy shocks affect young households mainly through labor income and old households mainly through asset returns. Most young households are hand-to-mouth and benefit from rising labor demand. Older households receive lower returns on their retirement savings. Almost all of the aggregate consumption response comes from working-age households. An unanticipated monetary easing redistributes welfare from the wealthiest old to the poorest young.

Keywords: HANK, Heterogeneous Agents, Life-Cycle Dynamics, Monetary Policy, Redistribution.

*bence.a.bardoczy@frb.gov

†mateo.velasquezgiraldo@frb.gov

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. Deriba Olana and 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.

Incorporating the life cycle of households into these macroeconomic models can improve their performance and credibility. First, abundant empirical evidence highlights age as an important determinant of households’ exposure to these macroeconomic policies and business cycles in general.¹ Second, since age is measured in most micro-level data sources, explicitly modeling it expands their testable predictions. Third, enhanced micro realism—by adding forces like the need to save for retirement—can alter the aggregate implications of these models and improve upon their known limitations.

In this paper, we embed a realistic life-cycle model of households into a heterogeneous agent New Keynesian framework, and find that it delivers on all three fronts. First, consistently with the empirical evidence, the model suggests that macroeconomic policy can affect households of different ages through dramatically different channels (Auclert, 2019). We study monetary policy shocks and find that young households are affected mostly by labor income and older households by asset returns. Second, in addition to matching income and wealth across the life cycle, our model generates predictions about the incidence of shocks across cohorts. For monetary policy shocks, 59% of the aggregate consumption response is due to households below the age of 40 and 97% is due to those below the age of 65; this is consistent with the available empirical evidence (Wong, 2019). Third, the inclusion of life-cycle dynamics at the micro level turns out to help the model 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).

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. The income process comes from Arellano et al. (2017) and has shocks that vary in their size, frequency, and persistence with age and income. As Janssens and McCrary (2023)

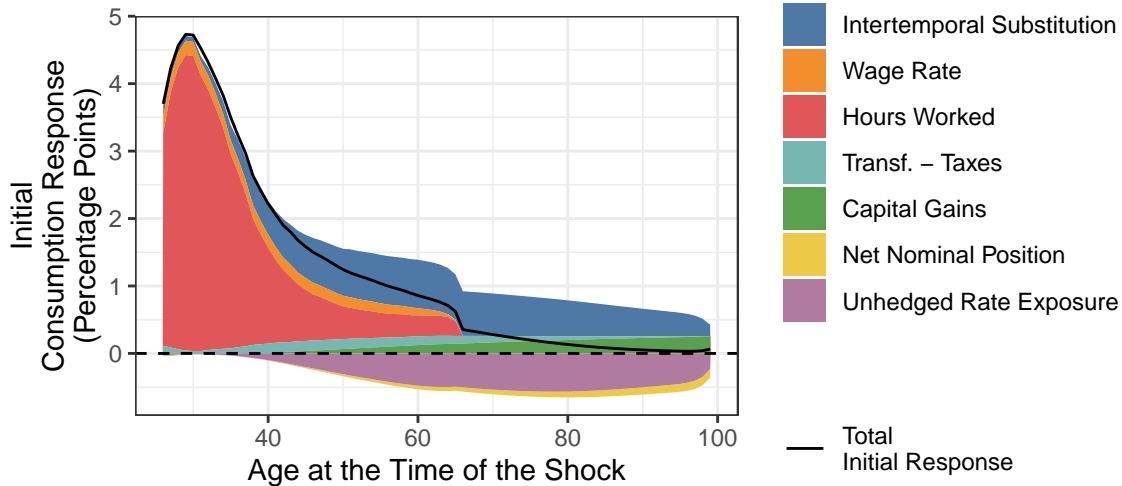
¹For example, Doepke and Schneider (2006), Adam and Zhu (2016), Greenwald et al. (2022), Pallotti et al. (2023), Fagereng et al. (2023) find heterogeneous exposure to inflation, asset prices, and interest rates across the life cycle. A large literature also documents that the exposure to labor-market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

point out, these features help to generate realistic distributions of wealth and MPCs. The luxury bequest motive is a form of non-homotheticity that raises the saving rates of wealthy and old households in line with the data (De Nardi, 2004). Borrowing constraints are key to elevating the MPCs, in particular for young households. As a result, consumption dynamics in our model reflect the main motives highlighted in the literature: households want to buffer against uncertainty, maintain their consumption during retirement, and leave bequests.

We embed the household block in a New Keynesian model that captures salient features of the business cycle. Wages and prices are set by monopolistic unions and firms subject to nominal rigidities and capital adjustment costs. Households can trade government bonds and firm equity through a financial intermediary. Wages and the price of firm equity rise moderately in response to an expansionary monetary policy shock, avoiding spurious redistribution between young workers and old firm owners (Broer et al., 2020). We match the unequal volatility of labor demand across the life cycle documented by Jaimovich et al. (2013). We model household portfolios as a flexible function of age and wealth, which we fit to the 2019 Survey of Consumer Finances (SCF). The monetary authority follows an inflation-targeting Taylor rule, and the fiscal authority adjusts a progressive tax function to stabilize government debt in the long run. The model approximates the responses of macroeconomic aggregates and asset prices to monetary policy shocks, and their unequal incidence across households.

Our model captures many channels of monetary transmission and redistribution that have been deemed important in the HANK literature. In particular, all potential channels of redistribution described by Auclert (2019) are operational and interact with heterogeneity along age and wealth. Figure 1 shows the contribution of these channels to the consumption responses to a monetary policy shock across the life cycle. Households aged 40 and below constitute less than 30% of the population but account for almost 60% of the initial response of aggregate consumption. Their strong response owes to their high MPCs and the disproportionate increase in their labor income. Asset returns play a more important role as households start to save for retirement in their forties. Capital gains and net nominal position capture the initial revaluation of real and nominal assets. Unhedged rate exposure captures the effect of persistently lower future interest rates on the affordability of consumption-saving plans. For retirees, the positive effect of intertemporal substitution and capital gains are largely offset by unhedged rate exposure. Overall, monetary policy affects the consumption of young households much more and through different channels than that of old households.

Short-run consumption responses paint an incomplete picture of the redistributive effects of monetary policy shocks. We compute a welfare measure that accounts for how the shock interacts with the uncertainty, constraints, and preferences of households. We find that younger cohorts gain and older cohorts lose, on average. 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 above 8% of their one-year consumption, while the gains of prime-age households in the



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 (dC_0) 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 change 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 $\{R_{t+s}^b, R_{t+s}^s\}_{s \geq 1}$. We solve for the effect of each individual channel non-linearly; therefore the sum of the mechanisms might slightly differ from the total.

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

lowest wealth quintile are almost 6% 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 in the two highest wealth quintiles above the age of 45. For the latter, the negative effect of asset returns can be as large as 10% of their one-year consumption in spite of the positive revaluation of their real assets at the time of the shock. The reason is that these household have very large ungedged interest rate exposure.

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), Stachurski and Toda (2019) identify 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, 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 wealth inequality 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), 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 assets; 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), Bullard et al. (2023), Beaudry et al. (2024) 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.² 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 (2006), Adam and Zhu (2016), Greenwald et al. (2022), Pallotti et al. (2023), Fagereng et al. (2023) are in this category. All of these papers find that age is a prominent dimension along

²Both 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.

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 an age-specific probability of death every year.³ 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 $l_{i,t}$ be hours worked. We adopt the following specification for income in working years ($\mathbf{a}_{i,t} \leq 65$)

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

³We use age-specific death probabilities from the SSA life-tables. We use cross-sectional probabilities from the year 2004. Figure C.3 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, $z_{i,t(\mathbf{a}_{i,t}=65)}$.⁴

We use the age-specific income shock process $z_{i,t}$ of Arellano et al. (2017) provided by Janssens and McCrary (2023).⁵ The key features of this shock process are nonlinear persistence and conditional heterogeneity of higher moments. Given the critical role of income shocks for the determination of savings and MPCs, we discuss the income process in greater detail 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 $l_{i,ss} = 1$, 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 $l_{i,t}$ must adjust to maintain

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

Out of the steady state, the distribution of $l_{i,t}$ across households is pinned down by an exogenous incidence function. It is a well-established fact that the hours of young workers fluctuate more over the business cycle than those of older workers (see, for example, Clark & Summers, 1981; Gomme et al., 2004; Jaimovich et al., 2013).⁶ 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 this evidence, we let labor demand for household i depend on both its own age and aggregate labor demand $l_{i,t} = \gamma(\mathbf{a}_{i,t}, L_t)$. Letting $\tilde{Y}_{\mathbf{a}} \equiv \int_{i:\mathbf{a}_i=\mathbf{a}} \tilde{y}_{i,t}$ be the total effective units per hour available from individuals of age \mathbf{a} , we set

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

⁴This 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), Kaplan and Violante (2014).

⁵We thank Eva Janssens for providing us with the optimally discretized income process.

⁶Sabelhaus and Song (2010), Guvenen et al. (2017) demonstrate a similar fact for fluctuations in earnings.

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, and displays the estimated incidence function.

2.3 Assets

Households save using two different assets: 1-period nominal deposits and firm equity. Deposits and equity have real return factors R_t^b and R_t^s , respectively. Because there is neither aggregate uncertainty nor portfolio adjustment costs, every agent expects that $R_t^b = R_t^s$ at all times. This implies that household portfolios are not determined by their optimizing behavior. However, household portfolios matter because unexpected aggregate shocks—“MIT shocks”—can generate differential returns ex-post.

The share of assets that household i invests in equity claims is a function of its age and its assets, $\zeta(\mathbf{a}_{i,t}, a_{i,t})$. We estimate the function $\zeta(\cdot, \cdot)$ using the 2019 Survey of Consumer Finances (SCF). Our measure of equity includes direct and indirect stock holdings and our measure of assets is total financial assets which importantly excludes housing. 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 by \mathcal{T} , that maps pretax into after-tax income. The function is

$$\mathcal{T}(y_{i,t}) = \lambda_t y_{i,t}^{1-\tau}, \quad (5)$$

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)⁷ and estimate the steady-state λ to match a 40% average tax rate on an income of \$300,000.⁸

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 $\mathcal{E} \equiv \int k_{i,t} dD^{\text{end}}(\mathbf{a}_{i,t} = 26)$. The endowments of newborns are heterogeneous and match the distribution of wealth for households aged between 21 and 25 years in the 2019 SCF. The government keeps the difference between the total assets of the dead and the total endowment of newborns.

⁷We use their estimate for pooled U.S. data and a narrow definition of transfers (Table 2, bottom row).

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

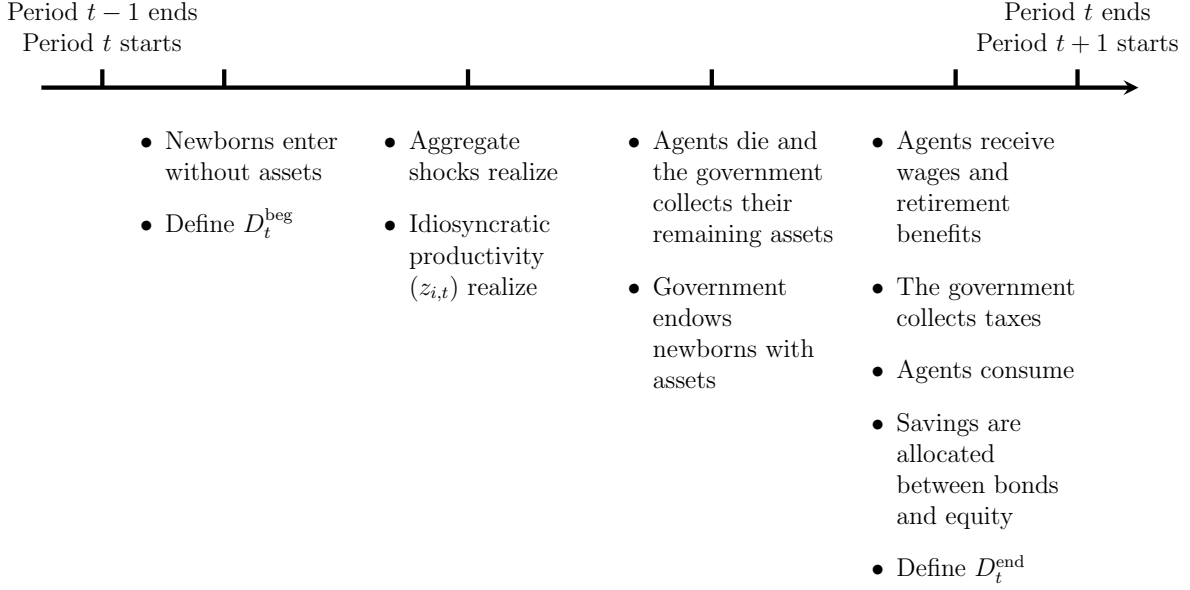


Figure 2: Summary of timing in the model

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 $\delta_{a_i,t}$ 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

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

where $1/\nu$ 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, newborns enter the model with $\mathbf{a} = 26$ and no initial assets.

Aggregate shocks occur and define the return to savings, while idiosyncratic productivity shocks also realize. Households die and receive utility from their bequests, which the government then collects while also handing newborns their endowment. Households then receive their wages, pension benefits, and transfers, and the government collects taxes. Finally, households decide how much to consume, and their remaining assets are distributed between government bonds and firm equity.

We now specify the recursive formulation of the optimization problem of a household that takes the sequences 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 $V_{\mathbf{a},t}(\alpha, z_t, a_{t-1})$ which is the value that the household expects at the beginning of the period, before knowing whether he will survive or not. The value is

$$\begin{aligned} V_{\mathbf{a},t}(\alpha, z_t, a_{t-1}) &= \delta_{\mathbf{a}_t} \phi(k_t) + \delta_{\mathbf{a}_t} \tilde{V}_{\mathbf{a},t}(\alpha, z_t, k_t) \\ \text{Where:} & \\ \tilde{R}_t &= R_t^b + \zeta(\mathbf{a}_{t-1}, a_{t-1}) \times (R_t^s - R_t^b) \\ k_t &= \tilde{R}_t \times a_{t-1}, \end{aligned} \tag{6}$$

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

Conditional on survival, the value function for a working-age household ($\mathbf{a} \leq 65$) is

$$\begin{aligned} \tilde{V}_{\mathbf{a},t}(\alpha, z_t, k_t) &= \max_{c_t} u(c_t) - v(l_t) + \beta \mathbb{E}_t[V_{\mathbf{a}+1,t+1}(\alpha, z_{t+1}, a_t)] \\ \text{subject to} & \\ m_t &= k_t + \mathcal{T}(w_t \times \tilde{y}_t \times l_t) \\ \tilde{y}_t &= \exp\{\alpha + f_{\mathbf{a}} + z_t\} \\ a_t &= m_t - c_t \\ a_t &\geq 0 \\ z_{t+1} &\sim \Pi_{\mathbf{a}_t}(z_t) \end{aligned} \tag{7}$$

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

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

where m_t denotes cash on hand.

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 $\tilde{R}_t = R_t^b + \zeta(99, a_{t-1}) \times (R_t^s - R_t^b)$.

We solve the household block using the method of endogenous gridpoints (Carroll, 2006). We use a 50-point doubly-nested exponential grid for end-of-period assets $a_{i,t}$, which runs from 10^{-4} to 500 times the average age-fixed-effect \bar{f}_a , and to which we add the constrained point $a_{i,t} = 0$ for a total of 51 points.

2.7 Calibration of the Household Block

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 the income process which requires a sequence of age-specific intercepts $\{f_a\}_{a=26}^{100}$ 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. The left panel of Figure 3 displays the empirical and model-implied percentiles of the income distribution, showing that our model replicates the age patterns of both its level and dispersion.

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 the age-varying percentiles implied by the model and those in the SCF.^{9,10} Our preference estimates, which we present in Table C.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 ($\kappa > 0$).¹¹ The right panel of Figure 3 compares the empirical and model-implied per-

⁹We use the standard Simulated Method of Moments loss function, with a diagonal weighting matrix that roughly rescales all moments to have similar magnitudes.

¹⁰We use the distribution of end-of-period assets $a_{i,t}$ in our model as the counterpart of SCF financial assets for this exercise.

¹¹See, for example, Carroll (1992), Attanasio et al. (1999), Gourinchas and Parker (2002), Cagetti (2003), De Nardi et al. (2010).



Our measure of income is the sum of wages and salaries, and Social Security and pension income from the 2019 SCF. Wealth-to-income ratios in the SCF are the ratio of financial assets to our measure of income. Our sample is households that report a strictly positive income and where the respondent is at least 21 years old. 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 income and wealth-to-income ratio percentiles for each age bin. Black dots correspond to the model-fitted counterparts to these percentiles.

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

centiles 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.

2.8 Wealth Distribution, MPCs, and Exposure to Interest Rates

The heterogeneous-agent macroeconomics literature has shown that the distributions of savings and MPCs across households in the economy modulate the effects of fiscal and monetary policy (Kaplan & Violante, 2014; Carroll et al., 2017; Auclert, 2019). 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. Additionally, we show that the life-cycle dynamics and saving motives in our model increase the heterogeneity in households' exposure to unexpected changes in interest rates. Studies like Auclert (2019), Beaudry et al. (2024) have highlighted this exposure as a mechanism of the distributive and aggregate effects of monetary policy.

Wealth distribution. The first feature that our model reproduces is the fact that wealth is more unequally distributed than income. The Gini coefficient of wealth in the model is

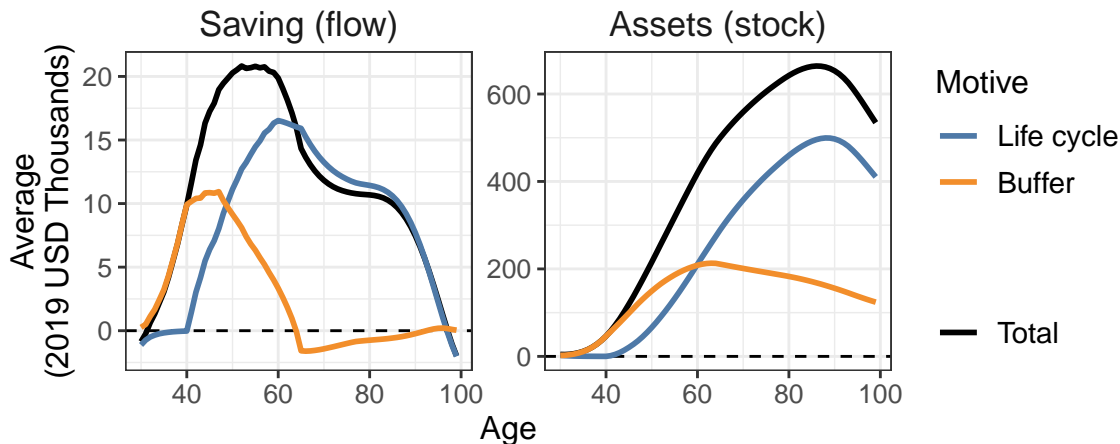
0.90, and in the SCF it is 0.86. In comparison, the model-implied Gini coefficient of income is 0.61 before taxes and 0.50 after taxes. In the SCF, the Gini coefficient of income (wages, salaries, social security, and pensions) is 0.5, not accounting for taxes. Reproducing such a large difference between income and wealth inequality has been difficult for studies featuring agents with identical preferences (Quadrini & Rios-Rull, 1997; Stachurski & Toda, 2019).¹² The success of our model in this dimension comes from explicitly representing 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 regarding the distribution of wealth (Huggett, 1996; De Nardi, 2004). It also comes from the particular discretization of the income process that we use which, as Janssens and McCrary (2023) show, matters for wealth inequality and MPCs.

Figure 4 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. 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.

Average MPCs. The second feature that we highlight is that our model generates MPCs of magnitudes that are consistent with empirical estimates. These estimates range roughly between 0.2 and 0.6 which is an order of magnitude greater than the MPCs predicted by standard representative-agent models.¹³ The left panel of Figure 5 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 wealth (total financial assets as a multiple of income, see Figure 3). Indeed, in their review of heterogeneous-agent models and MPCs, Kaplan and Violante (2022) conclude that one-asset models where agents

¹²Various studies reverse-engineer the income process or allow for preference heterogeneity to match wealth inequality (Castañeda et al., 2003; Carroll et al., 2017).

¹³See Jappelli and Pistaferri (2010), Carroll et al. (2017), Crawley and Theloudis (2024) for summaries of this literature and Fagereng et al. (2021) for estimates using Norwegian administrative data.



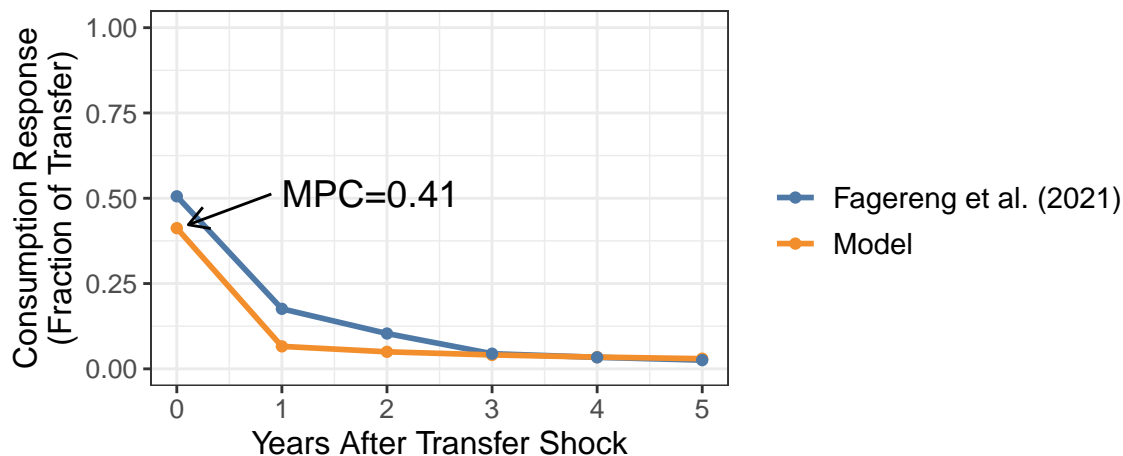
This figure decomposes the saving flows and accumulated asset stocks into the mechanisms that generate them. The decomposition follows Gourinchas and Parker (2002): “total” is our baseline model, “life-cycle” is a model in which income uncertainty is eliminated—each household receives the age-specific average post-tax income associated with its fixed effect (α)—and “buffer” is the difference between the two models. Because the “buffer” line is obtained as a residual, it contains the complementary effects between life-cycle and buffer-stock saving; for example, the fact that households get bequest utility from their buffer savings. This is why buffer assets do not approach 0 at the terminal age.

Figure 4: The Role of Different Mechanisms for Wealth Accumulation

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 feature of our model is that it generates substantial heterogeneity in MPCs without relying on heterogeneous or behavioral preferences; instead, the variation is mainly driven by wealth and age. Figure 6 depicts the average annual MPC for households 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 households 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), McKay and Wolf (2023).

The heterogeneous MPCs in Figure 6 co-vary with age and wealth: they fall 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 households from running down their assets: only around half 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), Bullard et al.



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

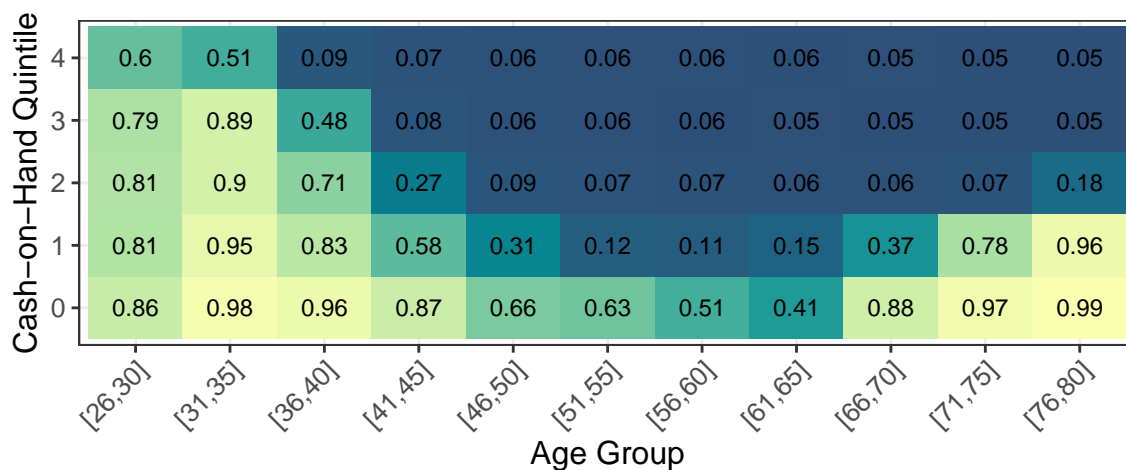
Figure 5: Intertemporal Marginal Propensities to Consume (iMPCs).

(2023), and also replicates the fact that a fraction of the elderly keep large stocks of savings (De Nardi et al., 2010).¹⁴

Figure 7 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 3 and that of alternative specifications in Figure C.4 of Appendix C.

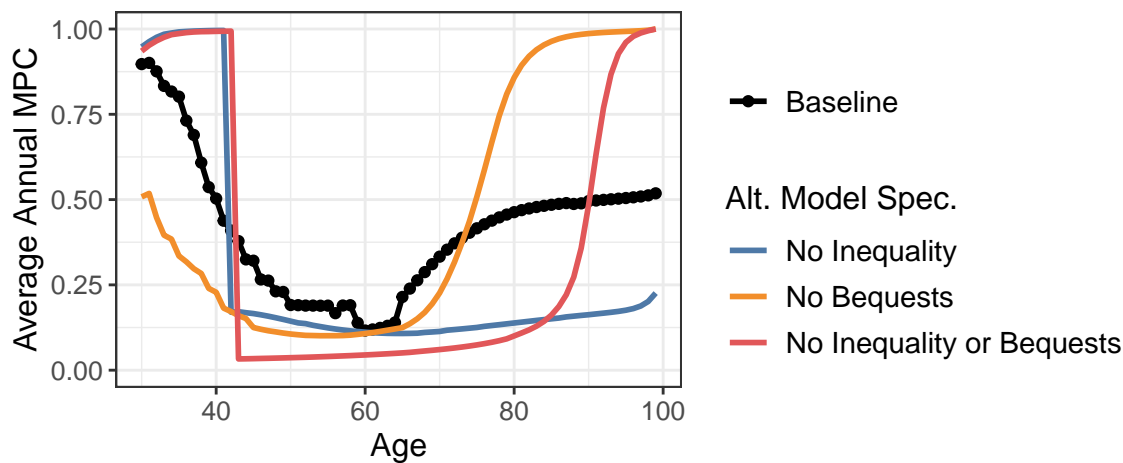
In our baseline model, MPCs fall gradually by age, and rise only moderately in old age. Without income and wealth inequality within age groups, young households quickly run down their inheritance and become hand-to-mouth. In this version of the model, the pressure to save for retirement reaches a critical 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.

¹⁴Bequests 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.



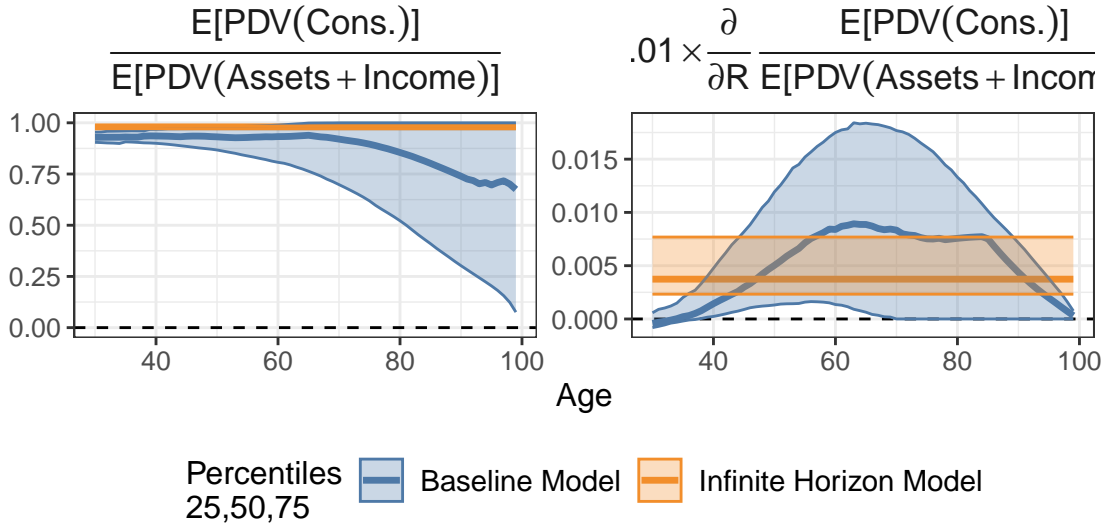
The figure depicts annual MPCs across the steady-state distribution of households in our model. Starting with the steady-state distribution, we group agents into the depicted age bins. For each age bin, we group agents into quintiles of their cash-on-hand ($m_{i,t}$). For each age and cash-on-hand group, 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 6: MPCs Across the Life Cycle and the Wealth Distribution.



The figure presents the average 1-year marginal propensity to consume across all agents of a given age in various specifications of the model. “Baseline” corresponds to our baseline model. “No bequests” removes utility from leaving bequests, keeping everything else unchanged. “No inequality” removes income uncertainty and inequality by giving every household the age-specific average post-tax income of the baseline model. “No inequality or bequests” removes utility from bequests from the “No inequality” model. Every model is estimated to match wealth-to-income ratios in the SCF 2019 and, therefore, they use different preference parameters.

Figure 7: MPCs in Alternative Models that Target Life-Cycle Wealth Ratios



The left panel depicts the ratio of the present value of future consumption to that of future income (after-tax wages and retirement benefits) plus current assets. The expectations and present-values are taken from the start of a given period, before the returns of that period have been applied. The right panel depicts the derivative of this ratio with respect to the equilibrium interest rate that applies to every period and asset, re-scaled to approximate a 1 p.p. fall in rates. Derivatives do not account for behavioral responses (consumption plans do not change). See the main text for a description of the infinite-horizon model.

Figure 8: The Cost of Consumption Plans and the Effect of Interest Rates

Exposure of consumption plans to interest rates. Households smooth their consumption across both time and shock realizations: their plans include periods and states in which consumption and income differ. These differences make the present values of their consumption and that of their lifetime resources differentially sensitive to interest rate changes. Because of this mismatch in sensitivities, households can have unhedged exposure to interest rate changes (Auclert, 2019). We explore the features of this mismatch in our model. In this analysis, we compare our life cycle model with an infinite horizon version in which we turn off the bequest motive, calibrate the (constant) survival probability to generate the same life expectancy as that of a 26-year-old in our baseline model, and the discount factor β to produce the same aggregate ratio of end-of-period assets to post-tax income; we also fix the income process on the grid and transition probabilities of age 45 in our baseline model.¹⁵

Because of voluntary bequests, not all households plan to consume all of their lifetime resources. The left panel of Figure 8 compares the ratio of the expected present-discounted value of consumption to that of income and assets, depicting its distribution across households of every age. In models without voluntary bequests, this ratio is 1.0 for every household because they expect to consume all of their lifetime resources. In our model, since unconsumed resources generate utility through bequests, the ratio can fall below 1.0. The figure shows how much variation the “luxury bequests” specification generates in households’

¹⁵The other preference parameters, income fixed effects, interest rates, and the taxation function are unchanged.

planned savings. By age 80, more than one quarter of households has decided not to leave bequests, while another quarter plans to consume less than 52% of their remaining lifetime resources.

We now illustrate how interest rates change the value of lifetime consumption relative to lifetime resources. The metric that we use is the derivative of the ratio between the value of future consumption to that of remaining lifetime resources with respect to the equilibrium interest rate,

$$\frac{\partial}{\partial R} \left\{ \frac{\mathbb{E}_t [\text{PDV}(\{c_{i,t+s}\}_{s=0}^{\infty})]}{\mathbb{E}_t [a_{i,t-1} + \text{PDV}(\{\text{Income}_{i,t+s}\}_{s=0}^{\infty})]} \right\}, \quad (9)$$

where “Income” denotes after-tax labor income and retirement benefits, and R the interest rate that applies to all periods and assets in equilibrium. The derivative does not account for behavioral responses to interest rate changes: it measures changes in the valuation of the original consumption plan and income path. A negative value of this derivative—as we find for most households in most periods—implies that a marginal reduction of interest rates makes the original consumption plan of the household less affordable relative to their lifetime resources. Such a change would induce the household to reduce its consumption, leave a smaller bequest, or both.

We find that interest rate reductions increase the cost of lifetime consumption relative to resources for most households, and that this effect is greater in our calibrated life-cycle model than in the infinite-horizon model. The right panel of Figure 8 displays the distribution of our measure from Equation 9, multiplied by -0.01 to approximate the effect of a 1 percentage-point reduction in interest rates. Our measure of sensitivity follows a hump shape across the life cycle, starting and ending close to 0 and peaking around the age of retirement. At age 65, the effect of the decline in interest rates on the median household would be to increase the value of its consumption relative to its resources by 0.9 percentage points. This is greater than the median 0.4 percentage point effect on its infinite-horizon counterpart. The same is true for a large part of the life cycle: the median effect is above its infinite-horizon counterpart from age 47 to 91. Even though households’ “human wealth” (Summers, 1981) appreciates, so does their planned consumption and the latter effect gains relative importance when households finance retirement consumption and bequests using savings that yield lower returns (Auclert et al., 2018; Beaudry et al., 2024).

Despite the lack of life-cycle motivations to save, households in the infinite-horizon model are also exposed to interest rates: reductions also appreciate their consumption relative to their lifetime resources. There are two main reasons behind this fact. The first is that households receive interest income from the buffer of savings that they maintain to insure against shocks. Lower rates reduce this stream of income, raising the relative value of planned consumption. The second reason is that, because the consumption function is concave, positive income shocks are smoothed over longer horizons than negative income shocks of the same size. This increases the expected duration of consumption relative to that of income. The life-cycle model features the same forces, but the additional need to save for retirement and bequests leads to greater heterogeneity.

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 aggregate consumption operates through the covariance of agents’ MPCs with various features of their income, consumption plans, and asset holdings.¹⁶ However, previous studies of monetary policy in heterogeneous-agent economies have abstracted away either from life-cycle considerations or from heterogeneity within age groups.¹⁷ Models that incorporate only life-cycle heterogeneity tend to have unrealistic implications for wealth or MPCs, or both (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.

3 A Life-Cycle HANK Model

Next, we embed our household model into a general equilibrium framework. 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), Kaplan et al. (2018), Alves et al. (2020).

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 $\{C_t, A_t^s, A_t^b, T_t, \Lambda_t, v'(L_t^*), u'(C_t^*)\}_{t=0}^\infty$. The inputs are the real wage w_t , labor demand L_t , return on nominal deposits R_t^b , return on stocks R_t^s , and the intercept of the retention function λ_t . The outputs are consumption C_t , savings in stocks A_t^s , savings in bonds A_t^b , taxes net of transfers T_t , bequests Λ_t , and the average costs and benefits of labor supply $v'(L_t^*), u'(C_t^*)$.

To define these aggregate outputs, consider the distributions of agents over states at the beginning and end of period t : D_t^{beg} and D_t^{end} . We define

¹⁶Specifically, their incomes, net nominal positions, and unhedged interest rate exposures.

¹⁷See Kaplan et al. (2018), Auclert et al. (2020) for examples without a life-cycle component, and Braun and Ikeda (2021), Bielecki et al. (2022) for examples with life-cycle differences but homogeneous cohorts.

$$\Lambda_t = \int \delta_{i,t} k_{i,t} dD_t^{\text{beg}} \quad T_t = \int y_{i,t} - \mathcal{T}(y_{i,t}) dD_t^{\text{end}}, \quad C_t = \int c_{i,t} dD_t^{\text{end}} \quad (10)$$

$$A_t^s = \int \zeta(\mathbf{a}_{i,t}, a_{i,t}) \times a_{i,t} dD_t^{\text{end}}, \quad A_t^b = \int (1 - \zeta(\mathbf{a}_{i,t}, a_{i,t})) \times a_{i,t} dD_t^{\text{end}}, \quad (11)$$

where income is

$$y_{i,t} = \begin{cases} w_t \times \tilde{y}_{i,t} \times l_{i,t} & \text{for } \mathbf{a}_{i,t} \leq 65, \\ d_{i,t} & \text{for } \mathbf{a}_{i,t} > 65. \end{cases} \quad (12)$$

We define the marginal utilities $v'(L_t^*)$, $u'(C_t^*)$ 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 welfare 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

$$\pi_t^w (\pi_t^w - 1) = \kappa_w \left(\frac{v'(L_t^*)}{u'(C_t^*)} - 1 \right) + \beta \mathbb{E}_t [\pi_{t+1}^w (\pi_{t+1}^w - 1)], \quad (13)$$

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 $v'(L_t^*)/u'(C_t^*)$ is a sufficient statistic that captures the impact of market power and distortionary taxes on labor supply. Its components are

$$v'(N_t^*) = \int L_t v'(l_{i,t}) \epsilon_w \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} dD_t^{\text{end}} (\mathbf{a}_{i,t} \leq 65), \quad (14)$$

$$u'(C_t^*) = \int u'(c_{i,t}) T'(y_{i,t}) y_{i,t} \left(\epsilon_w \frac{L_t}{n_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} - 1 \right) dD_t^{\text{end}} (\mathbf{a}_{i,t} \leq 65). \quad (15)$$

Financial intermediary. We introduce a financial intermediary to reconcile the portfolio of the household sector (taken from the SCF) with the net supply of assets in our model. The intermediary holds all assets in the economy. The difference between household wealth and asset supply is the net worth of the intermediary.

The balance sheet of the financial intermediary is

$$p_t + B_t = A_t^s + A_t^b + N_t, \quad (16)$$

where assets include stocks of total value p_t and bonds B_t , and liabilities include the stocks and bonds of the household sector (A_t^s and A_t^b) and the intermediary's own net worth N_t . Net worth evolves according to

$$N_t = R_t^s p_{t-1} + R_t^b B_{t-1} - R_t^s A_{t-1}^s - R_t^b A_{t-1}^b - d_t^{FI}, \quad (17)$$

where d_t^{FI} is a dividend paid to the owner of the intermediary. For simplicity, we assume that the owner is the government. The dividend payouts follow a rule

$$d_t^{FI} = \iota d^{FI} + (1 - \iota) [R_t^s p_{t-1} + R_t^b B_{t-1} - R_t^s A_{t-1}^s - R_t^b A_{t-1}^b - N] \quad (18)$$

with $\iota \in [0, 2 - R]$ is a smoothing parameter. Setting $\iota = 0$ implies that net worth is constant and the government budget is highly exposed to asset price fluctuations. Setting a higher ι dampens the impact of asset price fluctuations on the government budget.

Absent portfolio adjustment costs and aggregate uncertainty, stocks and bonds must offer the same expected return, implying that

$$R_t^e \equiv \mathbb{E}_t[R_{t+1}^b] = \mathbb{E} \left[\frac{R_t^n}{\pi_{t+1}} \right] = \mathbb{E}_t[R_{t+1}^s] = \mathbb{E}_t \left[\frac{d_{t+1} + p_{t+1}}{p_t} \right], \quad (19)$$

where R_t^e is the economy-wide real interest rate and R_t^n is the nominal interest rate set by the central bank. As we discussed earlier, the ex-post returns may differ

$$R_t^b = \frac{R_{t-1}^n}{\pi_t} \neq R_t^s = \frac{p_t + d_t}{p_{t-1}}. \quad (20)$$

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 $y_t = F(k_{t-1}, l_t) = \Theta k_{t-1}^\alpha l_t^{1-\alpha}$. 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

$$\pi_t(\pi_t - 1) = \kappa_p (\mu_p m c_t - 1) + \mathbb{E}_t \left[\frac{\pi_{t+1}(\pi_{t+1} - 1) Y_{t+1}}{R_t^e Y_t} \right], \quad (21)$$

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, $m c_t = w_t / F_L(K_{t-1}, L_t)$ is the real marginal cost, and Y_t is aggregate output.

The dynamics of investment is given by

$$Q_t = 1 + \psi \left(\frac{K_t}{K_{t-1}} - 1 \right), \quad (22)$$

$$R_t^e Q_t = \mathbb{E}_t \left[\alpha \frac{Y_{t+1}}{K_t} m c_{t+1} - \frac{I_{t+1}}{K_t} - \frac{\psi}{2} \left(\frac{K_{t+1}}{K_t} - 1 \right)^2 + \frac{K_{t+1}}{K_t} Q_{t+1} \right], \quad (23)$$

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 profit is

$$d_t = Y_t - w_t L_t - I_t - \frac{\psi}{2} \left(\frac{K_t}{K_{t-1}} - 1 \right)^2 K_{t-1} - \frac{\chi_p}{2} (\pi_t - 1)^2 Y_t. \quad (24)$$

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 the financial intermediary, and from running a progressive tax and transfer system. In sum, the primary surplus of the government is

$$S_t = T_t - G_t + \Lambda_t - \mathcal{E} + d_t^{FI}, \quad (25)$$

and its budget constraint is

$$B_t + S_t = R_t^b B_{t-1}. \quad (26)$$

Following Auclert et al. (2020), we assume that the government adjusts the intercept of the retention function λ_t according to the rule

$$\lambda_t = \lambda_{ss} - \phi \frac{B_{t-1} - B_{ss}}{Y_{ss}}. \quad (27)$$

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

$$R_t^n = R_{ss}^n + \phi_\pi (\pi_t - \pi_{ss}) + \varepsilon_t^{mp}, \quad (28)$$

where $\phi_\pi > 1$, and ε_t^{mp} is an exogenous monetary policy shock.

Equilibrium. Given a sequence of monetary policy and government spending shocks $\{\varepsilon_t^{mp}, G_t\}$, an exogenous distribution of endowments, and initial conditions D_{-1}, K_{-1}, B_{-1} , equilibrium is a sequence of allocations $\{Y_t, C_t, A_t, A_t^s, A_t^b, B_t, \Lambda_t, K_t, I_t, L_t, d_t, d_t^{FI}\}$ 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 financial intermediary, the government, and the central bank follow their policy rules;
- The balance sheet (16), no arbitrage (19), and realized return (20) conditions hold;
- Goods market clears

$$Y_t = C_t + G_t + I_t + \frac{\psi}{2} \left(\frac{K_t}{K_{t-1}} - 1 \right)^2 K_{t-1} + \frac{\chi_p}{2} (\pi_t - 1)^2 Y_t. \quad (29)$$

In Appendix B.3, we represent this economy as a directed acyclic graph (DAG).

3.2 Calibration

We calibrate the model in two steps. First, we calibrate the household block as described in sections 2.7 and 2.8. Second, we calibrate other blocks of the model, taking as given the household block and additional calibration targets from the literature and national statistics. Table C.2 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

$$\kappa_w = \frac{1}{1 + \Gamma_w} \frac{[1 - \beta(1 - \xi_w)]\xi_w}{1 - \xi_w}, \quad (30)$$

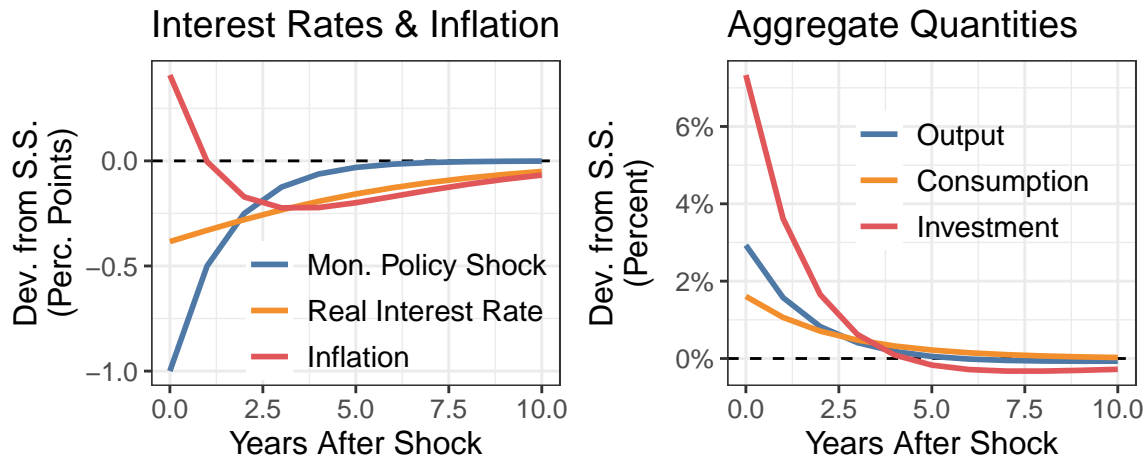
where β is the discount factor inherited from the household block, $\xi_w = 0.33$ is the annual frequency of wage adjustment from Grigsby et al. (2021). We set the real rigidity parameter $\Gamma_w = 5$ and the elasticity of substitution to the limit $\epsilon_w \rightarrow \infty$ following Auclert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

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 $K/Y = 2.23$. These are conventional choices. We calibrate the investment adjustment cost, ψ , to target a partial-equilibrium semi-elasticity of investment $d \log(I_t)/dR_t^e = -5$, 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 $\mu_p \rightarrow 1$ to shut down profits from monopoly power. Given this choice, equity price is equal to the value of capital, $p = K$.

Government. The household block pins households' savings in stocks A^s and bonds A_t^b . The firm block pins down the total value of stocks, p . We set government bonds to 46% of GDP. Together, these pin down the financial intermediary's net worth N . We set $\iota = 0.9$, implying smooth dividend payouts to the government. Government spending $G = 0.23$ is pinned down as a residual of the government budget. Following Auclert et al., 2020, we set the tax smoothing parameter to $\phi = 0.1$. 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



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

Figure 9: Aggregate Responses to an Expansionary Monetary Policy Shock

Taylor rule (Equation 28) is announced. The shock decays at a 0.5 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.

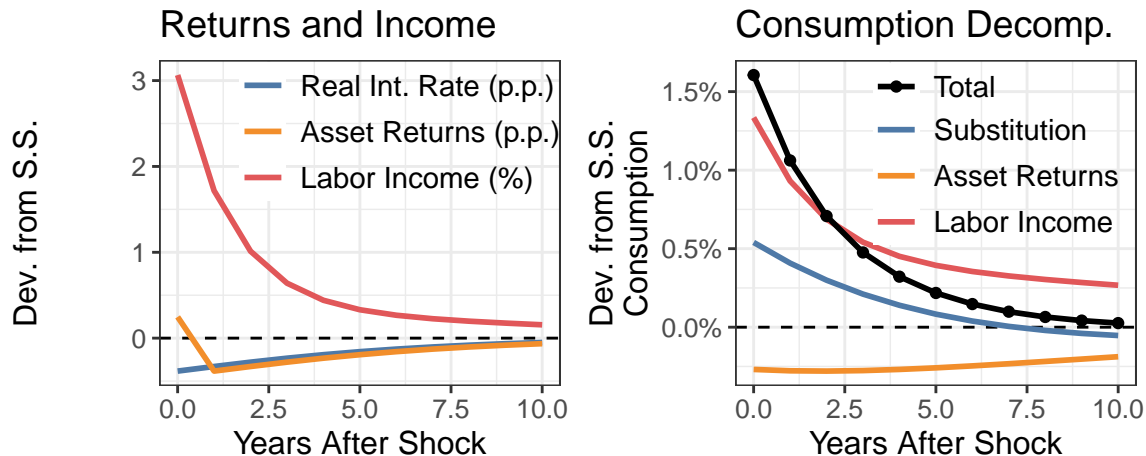
4.1 Aggregate Responses and Transmission Mechanisms

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 responses of macroeconomic aggregates to the shock are similar to those in typical New Keynesian models (for example, Christiano et al., 2005; Kaplan et al., 2018).¹⁸ The left panel of Figure 9 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 41 basis points and then decreases progressively, reaching its trough at 22 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 38 basis points on impact and remains below its steady-state level for the first 10 years. The right panel of Figure 9 shows the response of output, consumption, and investment, which all increase on impact. The response of investment is the greatest, at 7.3%, followed by output at 2.9% and consumption at 1.6%.

We now examine the drivers of the consumption response to the expansionary shock. We decompose the response into three components. First, the intertemporal substitution response to changes in expected interest rates, holding realized rates constant. Second, the response to changes in the realized rates of return on households’ assets, holding their

¹⁸Just 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 reliable 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 10: Decomposition of Aggregate Consumption Response to a Monetary Policy Shock

expected rates constant. And, third, changes in households’ labor income net of taxes and transfers. The left panel of Figure 10 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 on total assets initially increase due to a jump in the stock price; after the initial shock, realized asset returns are equal to the lagged ex-ante interest rate. The monetary policy shock produces a 3.1% increase in the labor income of households. Since our model features high MPCs, 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.

The right panel of Figure 10 decomposes the total response of consumption into the three main channels.¹⁹ The labor income channel accounts for a 1.3% increase in consumption on impact, out of 1.6% in total, and it continues to be the main driver of consumption as time passes. Intertemporal substitution generates only a 0.5% increase on impact. Persistently lower asset returns cause a 0.3% decrease in consumption that is also highly persistent.

Figure C.1 disaggregates the labor income channel into hours, wages, and tax rates. The wage Phillips curve is flat ($\kappa_w = 0.03$) and tax policy is inertial ($\phi = 0.1$), so the responses of the real wage and the tax rate are relatively small but persistent. Thus, the surge in hours drives most of the increase in labor income, especially in the short run. Hours increase more for younger workers who also have high MPCs. This is an example of the “matching multiplier” documented by Patterson (2023). All in all, higher hours account for more than

¹⁹In 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 at their steady state values.

two thirds of the total response.

Figure C.2 unpacks the asset returns channel. Stock and bond returns diverge on impact. Stocks are a real asset that appreciates in response to monetary easing. This is the capital gains channel. Bonds are a 1-period nominal asset whose return is eroded by surprise inflation. This is the net nominal exposure channel. Subsequently, R^b and R^s equalize and are low, consistent with monetary easing. Low expected returns affect the affordability of consumption-saving plans. This is the unhedged interest rate exposure channel. Net nominal exposure accounts for just a 4 basis points decline in consumption because surprise inflation itself is only 40 basis points in response to the monetary policy shock. Capital gains have a slightly larger impact, 9 basis points because the stock price jumps by more than 80 basis points. As Figure 4 shows, saving is highest among households between ages 40 and 80, who save for retirement and bequests. Their consumption-saving plans become less affordable when real returns fall, lowering the total consumption response on impact by 32 basis points.

Comparison to a HANK model without life cycle motives. A natural question is how the transmission of monetary policy to consumption differs between our model and a standard HANK model with perpetual youth households. Our explorations suggest that the answer to this question is sensitive to details such as the calibration of the income process. We think that one benefit of the life cycle model is that it facilitates a finer mapping to micro data that can inform these details.

The infinite horizon model that we use for our comparison is the same we used in Section 2.8 for the analysis of interest rate exposure. To recap, we turn off the bequest motive, calibrate the (constant) survival probability to generate the same life expectancy as that of a 26-year-old in our baseline model, and the discount factor β to produce the same aggregate ratio of end-of-period assets to post-tax income; we also fix the income process on the grid and transition probabilities of age 45 in our baseline model.

In our experiment, we take the general equilibrium responses of the inputs to the household block from the life cycle version of the model and feed them to the perpetual youth household block. This experiment isolates the differences in household behavior without confounding general equilibrium effects. Table 1 shows the decomposition of aggregate consumption in the life cycle and perpetual youth versions of the model.

There are three main takeaways from Table 1. First, the total consumption response is substantially larger in the infinite horizon model. The reason is that both models feature a high average MPC that makes labor income the dominant channel of transmission. The rise in labor income is driven mainly by labor demand and wages, which do not affect retired households. This dampens the labor income channel in the life cycle model. In addition, given our choice to target the same wealth to income ratio, the infinite horizon model ends up with a slightly higher MPC (0.49 vs 0.41 in the life cycle model). It achieves this high MPC by vastly overstating wealth inequality, with a Gini coefficient of 0.97. A large share of wealth is held by super high-income households. By age 45, the highest income state becomes very

Table 1: Decomposition of Aggregate Consumption Response to a Monetary Policy Shock

	Life-Cycle Model		Inf. Horizon Model	
	Impact	Cumulative	Impact	Cumulative
Total Response	1.61%	4.67%	2.47%	6.71%
<i>Mechanisms</i>				
Substitution	0.54%	0.82%	0.75%	0.66%
Labor Income	1.33%	7.42%	2.09%	11.51%
Asset Returns	-0.27%	-3.57%	-0.37%	-5.46%
<i>Labor income sub-components</i>				
Wage	0.12%	2.15%	0.25%	3.63%
Labor Demand	1.11%	1.96%	1.65%	3.24%
Taxes and Transfers	0.11%	3.31%	0.19%	4.64%
<i>Asset returns sub-components</i>				
Net Nominal Position	-0.04%	-0.35%	-0.04%	-0.49%
Capital Gains	0.09%	0.78%	0.12%	1.73%
Unhedged Rate Exposure	-0.32%	-4.00%	-0.45%	-6.70%

Note: Impact is defined as the percent deviation from steady state in the first year: dC_0/C_{ss} . Cumulative is defined as the present value of percent deviations from steady state: $\sum_{t=0}^{\infty} R_{ss}^{-t} dC_t/C_{ss}$.

persistent.²⁰ Without retirement or age-dependent mortality, these households accumulate very high levels of wealth, while many other households are hand-to-mouth. This highlights the ways in which subtle details can alter important aggregate properties of this class of models.

Second, the ranking of the intertemporal substitution channel between the two models is ambiguous. On the one hand, a higher average MPC tends to dampen the intertemporal substitution channel.²¹ On the other hand, a shorter life expectancy also dampens the intertemporal substitution channel. These two forces push in opposite directions. In our particular calibration, intertemporal substitution is stronger for the infinite horizon model on impact and for the life cycle model in the long run.

Third, the asset returns channels are stronger in the infinite horizon model. To illustrate the possible sources of the differences between the two models, consider the following decomposition

$$dC_0^{\text{NNP}} = \text{MPC} \times A_{ss}^b dR_0^b + \text{cov}(\text{MPC}_i, a_i^b) dR_0^b, \quad (31)$$

which holds for surviving households at the time of the shock.²² This is a standard result in a broad class of consumption-saving models (see Auclert, 2019). The aggregate consumption response equals the average MPC times the aggregate excess return plus the covariance between individual MPCs and bond holdings times the shock. The first term is larger in the infinite horizon model due to its larger average MPC. The covariance is negative and also larger in magnitude in the infinite horizon model due to its greater wealth inequality. On impact, these two differences roughly cancel out. Similar decompositions hold for capital gains and URE and, in general, the higher MPC and wealth in the infinite horizon model tend to dominate. Clearly, these results are sensitive to calibration details.

4.2 Monetary Policy Transmission Across the Life Cycle

We now explore how transmission and its channels vary across households of different ages. Young households are the main contributors to the increase in aggregate consumption and their response is due in great part to increased labor income. For older households, intertemporal substitution and asset returns play a more important role.


Figure 11 depicts the contribution of different cohorts to the change in aggregate consumption during the first year of the monetary shock. More than half of the response is due to households below the age of 40 and almost all is due to households below the age of retirement. 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%.²³ The right panel presents the

²⁰See Figure E.2 of Janssens and McCrary (2023).

²¹For example in the model of Auclert (2019), intertemporal substitution is proportional to $1 - \bar{\text{MPC}}$.

²²The equation does not hold for newborns because their endowment is not affected by the shock.

²³The 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.3 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 11: Incidence of the Initial Consumption Response Across Cohorts

cumulative shares; it shows that 59% of the initial response in consumption comes from households aged 40 or younger, and that 97% comes from working-age households (aged 65 or younger). The large contribution of young households to the aggregate consumption response and the steep decline of its incidence with age are consistent with the empirical evidence presented in studies like Wong (2019).²⁴

Table 2 decomposes the consumption response of cohorts aged 30, 50, and 70 both on impact and on the medium run. For working-age households, the labor income channel is the main driver of the consumption response. For older cohorts, the positive effect of intertemporal substitution and the negative effect of lower realized returns both increase in magnitude and partially offset each other. 30-year-old households increase their consumption by 4.72% in the first year, driven almost entirely by the increase in their labor income—the contribution of intertemporal substitution and asset returns is negligible. Most young households in our model are hand-to-mouth agents that simply consume the increased income that they receive: they have low savings and high MPCs (see Figures 3 and 6). For older households that have begun to accumulate assets for retirement and bequests, consumption responses become weaker and more mechanisms become relevant. For 50-year-olds, consumption increases by 1.24% on impact; labor income is still the channel with the greatest contribution (0.80%) but intertemporal substitution is almost as important (0.70%) and asset returns have a sizable negative effect (−0.27%). Finally, the consumption of 70-year-olds increases

demographic distribution of our simulated populations.

²⁴Wong (2019) finds that 83% of the 1-year consumption response to an interest rate shock is due to 25-34-year-olds, 15% is due to 35-64-year-olds, and only 2% is due to 65-75-year-olds. These are shares out of the total response of 25-75-year-olds. The analogous shares in our model are: 38% (25 to 34), 59% (35 to 64), and 3% (65 to 75).

Table 2: Decomposition of Consumption Response to a Monetary Policy Shock by Age

	On-impact, $s = 0$			Medium-run, $s = 5$		
	Age 30	Age 50	Age 70	Age 30	Age 50	Age 70
Total Response	4.72%	1.24%	0.28%	10.31%	4.39%	-0.09%
<i>Mechanisms</i>						
Substitution	0.07%	0.70%	0.63%	0.34%	2.11%	1.79%
Labor Income	4.64%	0.80%	0.09%	10.03%	3.86%	0.69%
Asset Returns	-0.01%	-0.27%	-0.44%	-0.08%	-1.59%	-2.57%
<i>Labor income sub-components</i>						
Wage	0.21%	0.16%	0.00%	1.92%	0.99%	0.00%
Labor Demand	4.37%	0.49%	0.00%	6.34%	1.84%	0.00%
Taxes and Transfers	0.03%	0.15%	0.09%	1.74%	1.02%	0.69%
<i>Asset returns sub-components</i>						
Net Nominal Position	0.00%	-0.03%	-0.07%	-0.02%	-0.15%	-0.33%
Capital Gains	0.00%	0.06%	0.16%	0.01%	0.31%	0.79%
Unhedged Rate Exposure	-0.01%	-0.31%	-0.54%	-0.08%	-1.76%	-3.01%

This table decomposes the consumption response of different cohorts to the monetary policy shock. Cohorts are indexed by their age at the time that the shock hits. All quantities are expressed as a percentage of a cohort’s consumption at time t in absence of the shock, $C_{t, \text{Age}}^{ss}$. The medium-run responses are present-discounted-values of cumulative consumption changes until five years after the shock: $\sum_{s=0}^5 (R^{ss})^{-s} dC_{t+s, \text{Age}+s}$. 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 the main text 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 $\{R_{t+s}^b, R_{t+s}^s\}_{s \geq 1}$. We solve for the effect of each individual channel non-linearly; therefore, the sum of the mechanisms can differ from the total.

by only 0.28% in the year of the shock, with the labor income channel contributing 0.09%, intertemporal substitution 0.63%, and asset returns -0.44% . The medium-run responses display similar patterns.

Rows five to seven of Table 2 decompose the income channel of the consumption response. The rise in hours is the largest contributor to the large consumption response of young households. For 30-year-old households, it generates a 4.37% increase in consumption, which is orders of magnitude greater than the on-impact contribution of wages and fiscal adjustments (transfers minus taxes). For 50-year-olds, these two channels gain relative importance: they generate consumption increases of 0.16% and 0.15% respectively, compared to the 0.49% increase caused by hours. The greater effect 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 6.15% at the time of the shock, while those of 50-year-olds increase only by 4.22%. Wages and hours do not mat-

ter for retirees, whose labor income response comes solely from the adjustment of taxes on their retirement benefits. Medium run responses show a greater role for taxes and transfers because of their slow and persistent response (see the left panel of Figure C.1).

Rows eight to ten of Table 2 decompose the effect of asset returns on the consumption of different cohorts into revaluations in their net nominal position, initial capital gains, and URE. The positive effect of capital gains from the on-impact revaluation of stocks generally outweighs the negative effects of the initial burst of inflation on bond returns (net nominal position): consumption would increase from these initial wealth effects. However, as pointed out by Auclert (2019), Beaudry et al. (2024), it would be wrong to conclude from this repricing that the new sequence of asset returns leaves these households better off. Figure 8 shows that there are households in our model with consumption plans that are highly sensitive to future rates of return. The persistent declines in returns can make these plans infeasible, leaving them worse off. These dynamic considerations, which indeed turn out to have a negative effect on consumption, are captured by the URE channel. The size of this effect increases with the age of the household at the time of the shock: it is -0.31% for 50-year-olds and -0.54% for 70-year-olds on-impact. This is enough to undo the positive wealth revaluation effects. Medium run responses show that the URE channel is also persistent and becomes the main driver of the consumption response of retirees in the five years that follow the shock.

4.3 The Redistributive Effects of Monetary Policy

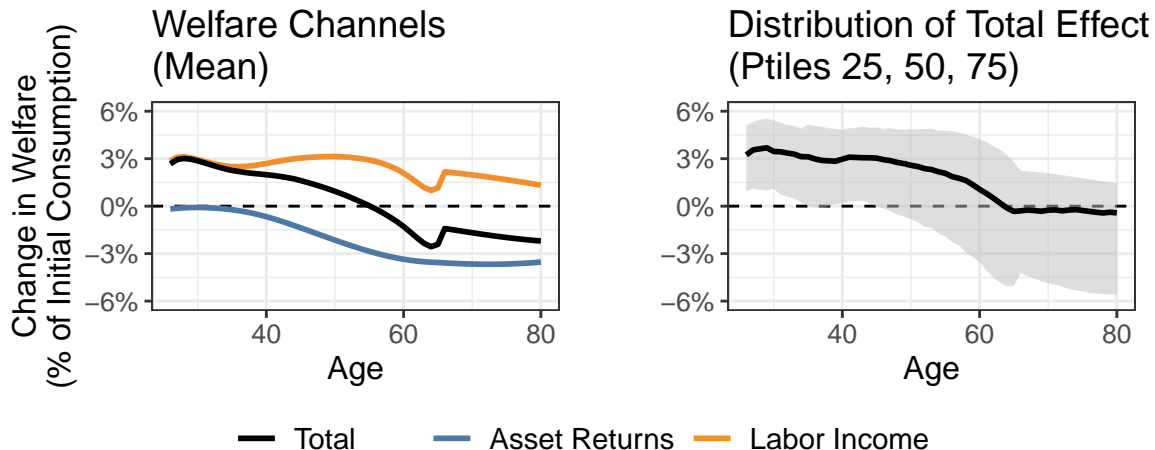
We now examine the redistributive effects of monetary policy. We use a welfare measure that captures the complex dynamic channels of redistribution. Importantly, our results should not be interpreted as prescriptions for optimal monetary policy. The study of systematic monetary policy raises considerations that are beyond the scope of this paper (for example, commitment). Nevertheless, our analysis elucidates who the winners and losers of an unanticipated monetary policy shock are and why.

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

$$\{w_{t+s}, L_{t+s}, R_{t+s}^b, R_{t+s}^s, \lambda_{t+s}\}_{s \geq 0}.$$

Using $V_{\mathbf{a}}^{SS}$ to denote 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 $V_{\mathbf{a}}^*$ its value function right after the announcement, the effect of the shock on the welfare of household i is

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



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 32). 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 12: Welfare Effects of a Monetary Policy Shock Across the Life Cycle

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:

$$\Delta_i \equiv \frac{1}{c_a^{SS}(\alpha_i, z_i, a_{i,t-1})} \times \frac{V_a^*(\alpha_i, z_i, a_{i,t-1}) - V_a^{SS}(\alpha_i, z_i, a_{i,t-1})}{\partial V_a^{SS}(\alpha_i, z_i, a_{i,t-1}) / \partial a_{i,t-1}}. \quad (32)$$

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.²⁵ 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 aged 55 and below, and negative for older households. The left panel of Figure 12 shows that the mean total effect is as high as 3% of a year's consumption for the youngest households, and that it declines steeply with age, generating losses greater than -2% for the oldest households.

The positive total welfare effect combines a labor income effect that is positive 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

²⁵The true equivalent variation solves $V_a^*(\alpha_i, z_i, a_{i,t-1}) = V_a^{SS}(\alpha_i, z_i, a_{i,t-1} + x)$. We rely on the approximation $V_a^{SS}(\alpha_i, z_i, a_{i,t-1} + x) \approx V_a^{SS}(\alpha_i, z_i, a_{i,t-1}) + x \times \partial V_a^{SS}(\alpha_i, z_i, a_{i,t-1}) / \partial a_{i,t-1}$.

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 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 12 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 welfare effect that declines with age: it is equivalent to 2.9% of consumption for 30-year-olds and 1.3% for 80-year-olds. The effect of the asset-returns channel is negligible for young households but large and negative for older households: it is equivalent to -0.1% of consumption for 30-year-olds and -3.5% 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 is generally larger than cross-age variation. The right panel of Figure 12 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. For 30-year-olds, the 25th and 75th percentiles of the distribution of welfare effects are, respectively, 1.1% and 5.4% of consumption; for 80-year-olds they become -5.6% and 1.4%. In both cases, these ranges are wider than the life-cycle variation of the median welfare effect, which is 3.6% for 30-year-olds and 0.4% for 80-year-olds. By this measure, while life-cycle differences in the welfare effects of monetary policy are significant, within-age heterogeneity is a larger share of their variation.

We now examine these heterogeneous welfare effects within age groups, which the extant literature on the redistributive effects of monetary policy across generations has not accounted for. While age is an important and relevant dimension of heterogeneity, different households of a given age are in vastly different economic shape (see Figure 3). These differences have a material impact on our conclusions about important quantities such as MPCs and welfare. For example, the right panel of Figure 12 also shows that the welfare impact on retirees is skewed. The welfare of the median retiree barely changes and the negative average effect shown in the left panel—and highlighted in previous studies—is generated by a reduced group of households that incur large losses; we will now characterize this group and the source of their losses.

The top panel of Figure 13 displays the average welfare effect of the expansionary shock for households in different age bins and different quintiles of the bin-specific wealth distribution. 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 shock redistributes welfare from older and wealthier households to younger and poorer ones. The poorest working-age households perceive gains comparable to 4.6% to 6.4% of their consumption, while the losses of the wealthiest retirees are greater than 7.9% of theirs. These conclusions 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.

The middle and bottom panels of Figure 13 decompose the average welfare effects of the top panel into its labor-income and asset-returns channels.

The middle panel of Figure 13 presents the welfare effects of labor income changes. The rise in labor income has generally positive welfare impacts. These effects are largest for working-age households in the lowest wealth quintile, for whom they range between 4.7% and 6.4% of consumption and 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 1.3% and 2.2% of consumption and that are due solely to tax reductions.

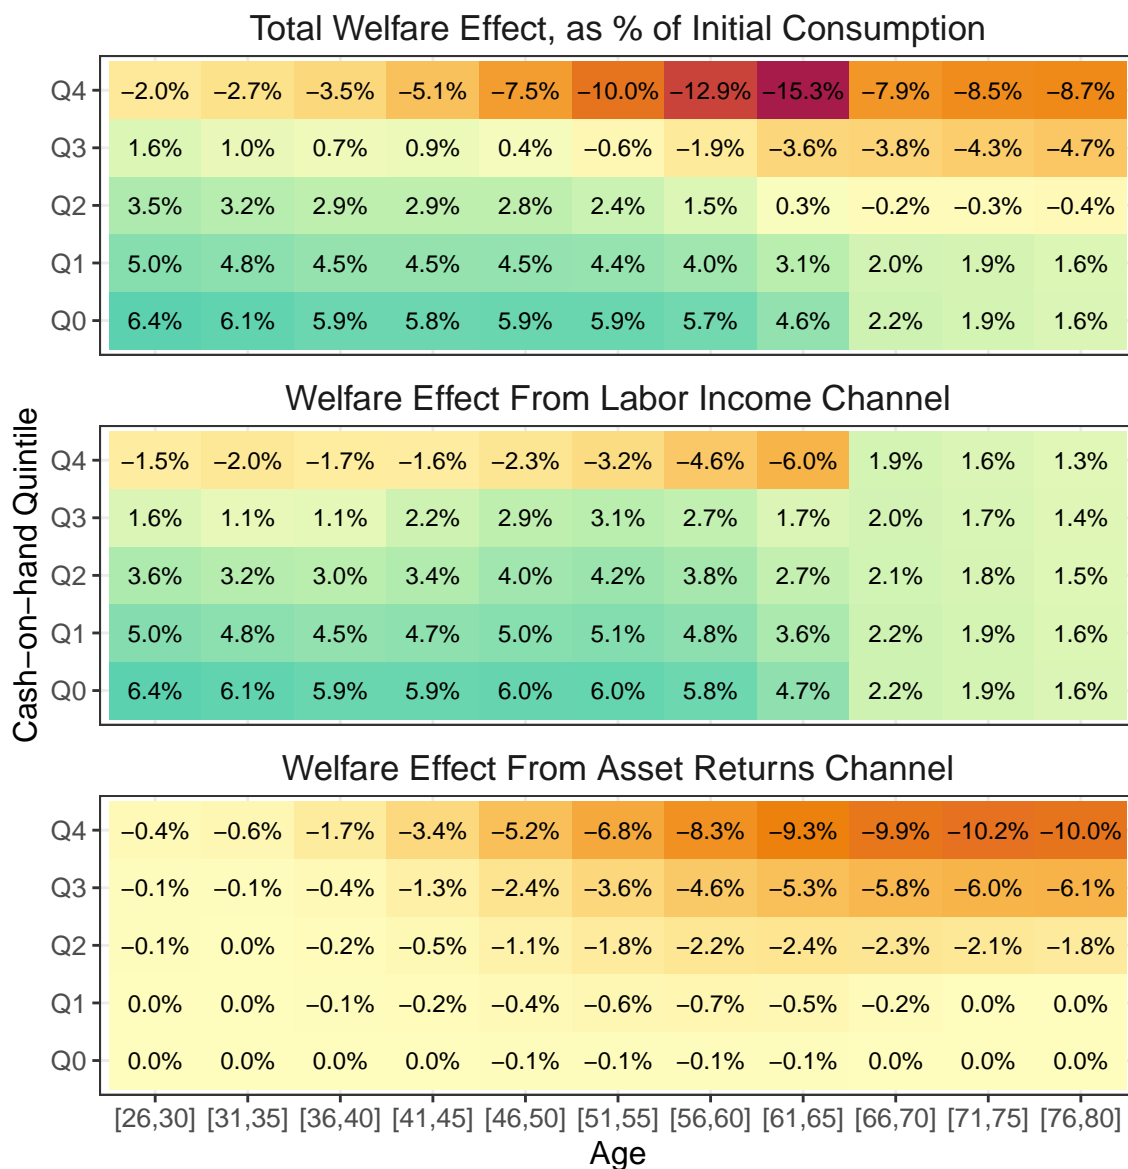
The bottom panel of Figure 13 presents the welfare effects of asset returns. The effects are negative at all ages and wealth quintiles. The magnitude of these effects is small for households in the first two wealth quintiles at every age. However, its magnitude increases with age and wealth, with the effects on the wealthiest retirees surpassing -10% of their consumption. These losses from the asset returns channel highlight two important points. First, the fact that they are the greatest for old and wealthy households—who also have the greatest exposure to the initial revaluation in stocks—serves to illustrate how initial wealth effects can differ from welfare effects. 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 on-impact aggregate consumption response was small; this is because the losses are concentrated in wealthy agents that smooth their consumption response over time.

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.²⁶ However, it affects the welfare implications of the model, since most households are never on their labor supply curve.²⁷ We maintain the assumption for lack of a better alternative, but note that it is the source of the negative welfare impacts from the labor channel on wealthy working households seen in the middle panel of Figure 13.

As a robustness check, Figure C.6 reports the welfare effects 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 age of retirement. The general pattern of welfare effects

²⁶First, 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).

²⁷See Huo and Ríos-Rull (2020), Gerke et al. (2023) for detailed discussions of this issue of households “working against their will” in New Keynesian models.



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 32). 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 13: Welfare Effects of a Monetary Policy Shock by Age and Wealth

remains the same: young and poor households benefit the most, old and rich households benefit the least, or lose, from monetary easing.

5 Conclusions

We develop a heterogeneous agent New Keynesian model that explicitly represents the life cycle of households. Our model achieves empirically realistic levels of income and wealth inequality within and across cohorts and reproduces several key facts about MPCs: they are high on average and decline with wealth and age until retirement (Fagereng et al., 2021; Colarieti et al., 2024). Matching these moments simultaneously has been challenging for existing models, but is achieved naturally in ours as a by-product of its life-cycle calibration. These features make our model a good laboratory to study the aggregate and distributional effects of macroeconomic policies and shocks.

We apply the model to study the transmission mechanisms and redistributive effects of a monetary policy shock. At the aggregate level, our model reproduces the key findings of the HANK literature. The response of consumption is similar to that in standard representative agent New Keynesian models but is driven to a large extent by a general equilibrium increase in labor demand, as opposed to pure intertemporal substitution (Kaplan et al., 2018). Our model allows us to refine these predictions along several dimensions. First, more than half of the total consumption response is due to households below the age of 40 and almost all of it is due to households below the age of retirement. This prediction is consistent with the evidence of Wong (2019). Second, labor income effects are indeed the key driver of consumption for young households who have high MPCs and are disproportionately exposed to fluctuations in labor demand. However, for older households who rely more on their savings to finance consumption, asset returns play a much greater role. Third, the negative income effect from persistently lower interest rates can overwhelm the positive wealth effect from the initial surge in stock prices. This explains why an unanticipated expansionary shock hurts older and wealthier households and benefits younger and poorer households. All in all, age emerges as an important dimension of heterogeneity.

Future work could extend our model in important ways. First, our model only features financial assets and abstracts away from non-financial assets including housing and durables as well as all forms of debt. A more extensive representation of household balance sheets would better approximate their exposure to changes in asset returns. Second, retirement is exogenous and features no idiosyncratic uncertainty. The inclusion of, for example, health shocks, could meaningfully alter the consumption-saving behavior of retirees.

Our model can be applied to many more policy-relevant questions. For example, the macroeconomic effects of fiscal consolidation and the impact of demographic transitions on the natural rate of interest. We look forward to exploring these questions in future work.

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. <https://doi.org/10.1111/jeea.12155>
- Aguiar, M. A., Bils, M., & Boar, C. (2020, January). *Who Are the Hand-to-Mouth?* <https://doi.org/10.3386/w26643>
- 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. <https://doi.org/10.1111/jmcb.12761>
- 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. <https://doi.org/10.1086/706686>
- Arellano, M., Blundell, R., & Bonhomme, S. (2017). Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. *Econometrica*, 85(3), 693–734. <https://doi.org/10.3982/ECTA13795>
- Attanasio, O. P., Banks, J., Meghir, C., & Weber, G. (1999). Humps and Bumps in Lifetime Consumption. *Journal of Business & Economic Statistics*, 17(1)jstor 1392236, 22–35. <https://doi.org/10.2307/1392236>
- Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6), 2333–2367. <https://doi.org/10.1257/aer.20160137>
- 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. https://doi.org/10.1162/rest_a.01072
- 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. <https://doi.org/10.3982/ECTA17434>
- Auclert, A., Rognlie, M., & Straub, L. (2018, September). *The Intertemporal Keynesian Cross*. <https://doi.org/10.3386/w25020>
- Auclert, A., Rognlie, M., & Straub, L. (2020, January). *Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model*. <https://doi.org/10.3386/w26647>
- Beaudry, P., Cavallino, P., & Willems, T. (2024, May). *Life-cycle Forces make Monetary Policy Transmission Wealth-centric*. <https://doi.org/10.3386/w32511>
- Benhabib, J., & Bisin, A. (2018). Skewed Wealth Distributions: Theory and Empirics. *Journal of Economic Literature*, 56(4), 1261–1291. <https://doi.org/10.1257/jel.20161390>
- 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. <https://doi.org/10.1257/aer.20151684>
- Bielecki, M., Brzoza-Brzezina, M., & Kolasa, M. (2022). Intergenerational Redistributive Effects of Monetary Policy. *Journal of the European Economic Association*, 20(2), 549–580. <https://doi.org/10.1093/jeea/jvab032>

- 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 2021-20a). Federal Reserve Bank of Atlanta. Retrieved February 29, 2024, from <https://econpapers.repec.org/paper/fipfedawp/93475.htm>
- 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. <https://doi.org/10.1093/restud/rdy060>
- Bullard, J. B., DiCecio, R., Singh, A., & Suda, J. (2023, July 19). *Optimal Macroeconomic Policies in a Heterogeneous World*. <https://doi.org/10.2139/ssrn.4633791>
- Cagetti, M. (2003). Wealth Accumulation Over the Life Cycle and Precautionary Savings. *Journal of Business & Economic Statistics*, 21(3), 339–353. <https://doi.org/10.1198/073500103288619007>
- Carroll, C. D. (1992). The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, 1992(2)jstor 2534582, 61–156. <https://doi.org/10.2307/2534582>
- Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis*. *The Quarterly Journal of Economics*, 112(1), 1–55. <https://doi.org/10.1162/003355397555109>
- 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. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3), 312–320. <https://doi.org/10.1016/j.econlet.2005.09.013>
- 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. <https://doi.org/10.1257/mac.20180286>
- 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), <https://onlinelibrary.wiley.com/doi/pdf/10.3982/QE694>, 977–1020. <https://doi.org/10.3982/QE694>
- 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. <https://doi.org/10.1086/375382>
- 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)jstor 10.1086/426038, 1–45. <https://doi.org/10.1086/426038>
- Clark, K. B., & Summers, L. H. (1981). Demographic Differences in Cyclical Employment Variation. *The Journal of Human Resources*, 16(1)jstor 145219, 61–79. <https://doi.org/10.2307/145219>

- Colarieti, R., Mei, P., & Stantcheva, S. (2024, March 4). *The How and Why of Household Reactions to Income Shocks* (w32191). National Bureau of Economic Research. <https://doi.org/10.3386/w32191>
- Crawley, E., & Theloudis, A. (2024, April 18). *Income Shocks and their Transmission into Consumption*. <https://doi.org/10.48550/arXiv.2404.12214>
- De Nardi, M. (2004). Wealth Inequality and Intergenerational Links. *The Review of Economic Studies*, *71*(3), 743–768. <https://doi.org/10.1111/j.1467-937X.2004.00302.x>
- De Nardi, M. (2015, April). *Quantitative Models of Wealth Inequality: A Survey*. <https://doi.org/10.3386/w21106>
- 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. <https://doi.org/10.1086/651674>
- Doepke, M., & Schneider, M. (2006). Inflation and the Redistribution of Nominal Wealth. *Journal of Political Economy*, *114*(6), 1069–1097. <https://doi.org/10.1086/508379>
- 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. <https://doi.org/10.1257/mac.20190211>
- 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*. <https://doi.org/10.2139/ssrn.4744547>
- Gomme, P., Rogerson, R., Rupert, P., & Wright, R. (2004). The Business Cycle and the Life Cycle. *NBER Macroeconomics Annual*, *19*, 415–461. <https://doi.org/10.1086/ma.19.3585347>
- Gourinchas, P.-O., & Parker, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, *70*(1), 47–89. <https://doi.org/10.1111/1468-0262.00269>
- Greenwald, D., Leombroni, M., Lustig, H. N., & Van Nieuwerburgh, S. (2022, July 12). *Financial and Total Wealth Inequality with Declining Interest Rates*. <https://doi.org/10.2139/ssrn.3789220>
- Grigsby, J., Hurst, E., & Yildirmaz, A. (2021). Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data. *American Economic Review*, *111*(2), 428–471. <https://doi.org/10.1257/aer.20190318>
- Guerrieri, V., & Lorenzoni, G. (2017). Credit Crises, Precautionary Savings, and the Liquidity Trap*. *The Quarterly Journal of Economics*, *132*(3), 1427–1467. <https://doi.org/10.1093/qje/qjx005>
- 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. <https://doi.org/10.1257/aer.p20171094>

- 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. <https://doi.org/10.3386/w29775>
- Heathcote, J., Storesletten, K., & Violante, G. L. (2017). Optimal Tax Progressivity: An Analytical Framework*. *The Quarterly Journal of Economics*, *132*(4), 1693–1754. <https://doi.org/10.1093/qje/qjx018>
- Huggett, M. (1996). Wealth distribution in life-cycle economies. *Journal of Monetary Economics*, *38*(3), 469–494. [https://doi.org/10.1016/S0304-3932\(96\)01291-3](https://doi.org/10.1016/S0304-3932(96)01291-3)
- Huo, Z., & Ríos-Rull, J.-V. (2020). Sticky Wage Models and Labor Supply Constraints. *American Economic Journal: Macroeconomics*, *12*(3), 284–318. <https://doi.org/10.1257/mac.20180290>
- 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. <https://doi.org/10.1257/aer.103.7.3022>
- Janssens, E. F., & McCrary, S. (2023, May 17). *Finite-State Markov-Chain Approximations: A Hidden Markov Approach*. <https://doi.org/10.2139/ssrn.4137592>
- Jappelli, T., & Pistaferri, L. (2010). The Consumption Response to Income Changes. *Annual Review of Economics*, *2*(1), 479–506. <https://doi.org/10.1146/annurev.economics.050708.142933>
- Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, *108*(3), 697–743. <https://doi.org/10.1257/aer.20160042>
- Kaplan, G., & Violante, G. L. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, *82*(4), 1199–1239. <https://doi.org/10.3982/ECTA10528>
- 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. <https://doi.org/10.1257/aer.20150063>
- McKay, A., & Wolf, C. K. (2023). Monetary Policy and Inequality. *Journal of Economic Perspectives*, *37*(1), 121–144. <https://doi.org/10.1257/jep.37.1.121>
- Nakamura, E., & Steinsson, J. (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models*. *The Quarterly Journal of Economics*, *123*(4), 1415–1464. <https://doi.org/10.1162/qjec.2008.123.4.1415>
- 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. <https://doi.org/10.3386/w31896>

- Patterson, C. (2023). The Matching Multiplier and the Amplification of Recessions. *American Economic Review*, 113(4), 982–1012.
<https://doi.org/10.1257/aer.20210254>
- Quadrini, V., & Rios-Rull, J.-V. (1997). Understanding the U.S. distribution of wealth. *Federal Reserve Bank of Minneapolis. Quarterly Review - Federal Reserve Bank of Minneapolis*, 21(2), 22–36. Retrieved March 3, 2024, from
<https://www.proquest.com/docview/227763405/abstract/6FFFC6967EB4845PQ/1>
- Rotemberg, J. J. (1982). Monopolistic Price Adjustment and Aggregate Output. *The Review of Economic Studies*, 49(4)jstor 2297284, 517–531.
<https://doi.org/10.2307/2297284>
- Sabelhaus, J., & Song, J. (2010). The great moderation in micro labor earnings. *Journal of Monetary Economics*, 57(4), 391–403.
<https://doi.org/10.1016/j.jmoneco.2010.04.003>
- 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. <https://doi.org/10.1016/j.jet.2019.04.001>
- Summers, L. H. (1981). Capital Taxation and Accumulation in a Life Cycle Growth Model. *The American Economic Review*, 71(4)jstor 1806179, 533–544. Retrieved August 12, 2024, from <https://www.jstor.org/stable/1806179>
- Wong, A. (2019). *Refinancing and The Transmission of Monetary Policy to Consumption*. <https://jfhoude.wiscweb.wisc.edu/wp-content/uploads/sites/769/2019/09/Arlene-Wong-refinancing.pdf>

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.

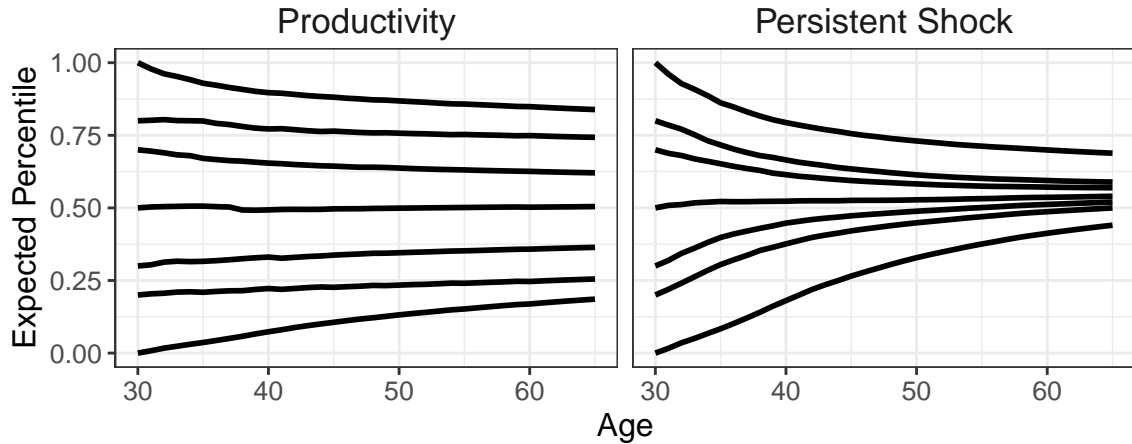
Productivity dynamics, especially their persistence, can have material effect on the behavior of consumption and savings. The optimal saving rate of a low-income household is very different depending on where it thinks its income will converge over time, at what rate, and with what degree of certainty. 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

Persistent Productivity Shock Grid



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

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.

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

$$\sigma(\ln(\gamma(\mathbf{a}_{i,t}, L_t))) \approx \sigma(\ln(L_t)) \times \left[\varepsilon_{\mathbf{a}_{i,t}} + 1 - \frac{\sum_a \tilde{Y}_a \varepsilon_a}{\sum_a \tilde{Y}_a} \right]. \quad (33)$$

Equation 33 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 33 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.

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

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 take a flexible functional form $\zeta(\cdot, \cdot; \vartheta)$ with parameters ϑ and estimate ϑ as

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

Our sample \mathcal{I} contains respondents in SCF-2019 who report strictly positive income and non-negative assets and who are above the age of 21. For the functional form of $\zeta(\cdot, \cdot; \vartheta)$, 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. We 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. 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 $\zeta(\cdot, \cdot)$. The estimated function captures various well-known empirical features, like a large fraction of the population not investing in equities and a steep relationship between equity share and wealth. To map into our model, we evaluate the estimated net $\zeta(\cdot, \cdot; \hat{\vartheta})$ on our asset grid at every age.




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

B Appendix to Section 3

This section contains the setup and detailed derivations of the New Keynesian model.

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 the 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.

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

$$y_{jt} = \left(\frac{p_{jt}}{P_t} \right)^{-\epsilon_p} Y_t \quad (34)$$

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

$$P_t = \left(\int_0^1 p_{jt}^{1-\epsilon_p} dj \right)^{\frac{1}{1-\epsilon_p}} \quad (35)$$

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

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

The FOC for l_{jt} is

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

Note that η_{jt} is the marginal profit from producing and selling an additional unit. The marginal profit equals the marginal revenue minus the marginal cost. So we can see from here that the real marginal cost is $mc_{jt} = w_t / \partial_L F(k_{jt-1}, l_{jt})$. In symmetric equilibrium,

$$mc_t = \frac{w_t}{\partial_L F(K_{t-1}, L_t)}. \quad (36)$$

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,

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

Combining the FOCs yields

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

Substituting $\eta_t = 1 - mc_t$ 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:

$$\pi_t (\pi_t - 1) = \underbrace{\frac{\epsilon_p}{\chi_p} \frac{\epsilon_p - 1}{\epsilon_p}}_{\kappa_p} \left(\frac{\epsilon_p}{\epsilon_p - 1} mc_t - 1 \right) + \mathbb{E}_t \left[\frac{\pi_{t+1} (\pi_{t+1} - 1) Y_{t+1}}{R_t^e} \frac{Y_{t+1}}{Y_t} \right]. \quad (37)$$

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 is

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

In symmetric equilibrium, we have that

$$\psi \left(\frac{K_t}{K_{t-1}} - 1 \right) = Q_t - p_t^I \quad (38)$$

$$R_t^e Q_t = mc_{t+1} \partial_k F_{t+1}(\cdot) + p_{t+1}^I (1 - \delta_k) - \Psi \left(\frac{K_{t+1}}{K_t} \right) + \frac{K_{t+1}}{K_t} (Q_{t+1} - p_{t+1}^I) \quad (39)$$

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 $\{Y_t, w_t, R_t^e, p_t^I\}$ and initial condition K_{-1} , the production block returns nine sequences of outputs $\{L_t, K_t, I_t, Q_t, mc_t, \pi_t, p_t, d_t, R_t^s\}$.

- **Production function.** Gives labor:

$$Y_t = F(K_{t-1}, L_t) = \Theta_t K_{t-1}^\alpha L_t^{1-\alpha} \quad (40)$$

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

$$mc_t = \frac{w_t}{F_L(K_{t-1}, L_t)} = \frac{1}{1-\alpha} \frac{w_t L_t}{Y_t} \quad (41)$$

- **Phillips curve.** Gives inflation:

$$\pi_t(\pi_t - 1) = \kappa_p \left(\frac{\epsilon_p}{\epsilon_p - 1} mc_t - 1 \right) + \mathbb{E}_t \left[\frac{\pi_{t+1}(\pi_{t+1} - 1) Y_{t+1}}{R_t^e Y_t} \right] \quad (42)$$

- **Marginal Q.** Gives Q_t :

$$Q_t = \psi \left(\frac{K_t}{K_{t-1}} - 1 \right) + p_t^I \quad (43)$$

- **Investment Euler.** Gives capital:

$$R_t^e Q_t = \mathbb{E}_t \left[\alpha \frac{Y_{t+1}}{K_t} mc_{t+1} - p_{t+1}^I \frac{I_{t+1}}{K_t} - \frac{\psi}{2} \left(\frac{K_{t+1}}{K_t} - 1 \right)^2 + \frac{K_{t+1}}{K_t} Q_{t+1} \right] \quad (44)$$

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

$$I_t = K_t - (1 - \delta_k) K_{t-1} \quad (45)$$

- **Profit.** Gives dividends:

$$d_t = Y_t - w_t L_t - p_t^I I_t - \Psi \left(\frac{K_t}{K_{t-1}} \right) K_{t-1} - \frac{\chi_p}{2} (\pi_t - 1)^2 Y_t \quad (46)$$

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

$$p_t = \mathbb{E}_t \left[\frac{d_{t+1} + p_{t+1}}{R_t^e} \right] \quad (47)$$

- **Return on equity.**

$$R_t^s = \frac{p_t + d_t}{p_{t-1}} \quad (48)$$

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. We 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 $\pi = 1$. Then, equations (43) and (42) imply immediately that

$$Q = 1, \quad mc = \frac{\epsilon_p - 1}{\epsilon_p}. \quad (49)$$

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

$$Y = \left(\frac{wL}{Y}\right)^{-1} wL, \quad K = \left(\frac{K}{Y}\right) Y \quad (50)$$

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

$$\alpha = 1 - \frac{wL}{Y} \frac{1}{mc}, \quad \delta_k = 1 + \alpha \frac{Y}{K} mc - R, \quad \Theta = \frac{Y}{K^\alpha L^{1-\alpha}}. \quad (51)$$

Finally, compute investment and flow profits (dividends)

$$I = \delta K, \quad d = Y - wL - I. \quad (52)$$

B.1.5 Block Jacobians

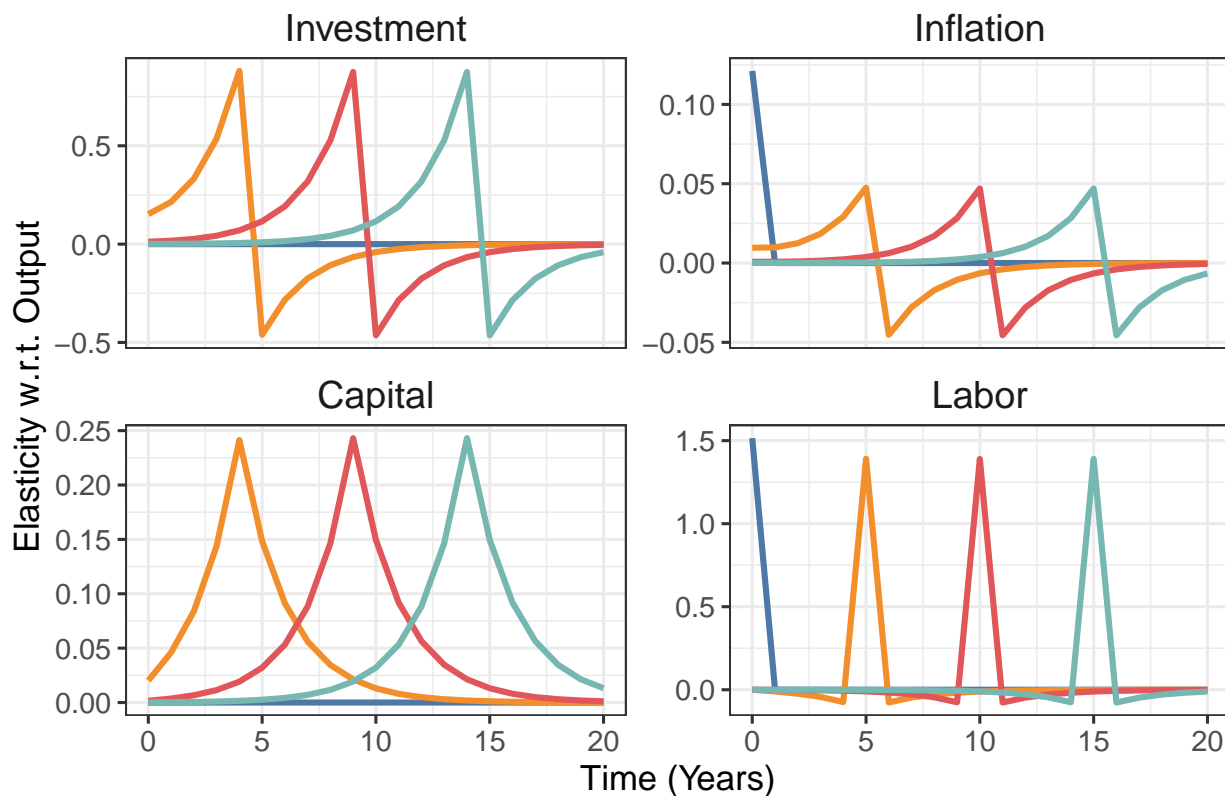


Figure B.1: Jacobians of the Firm Block with Respect to Output

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 $\partial \log O_t / \partial \log R_s^e$ 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.

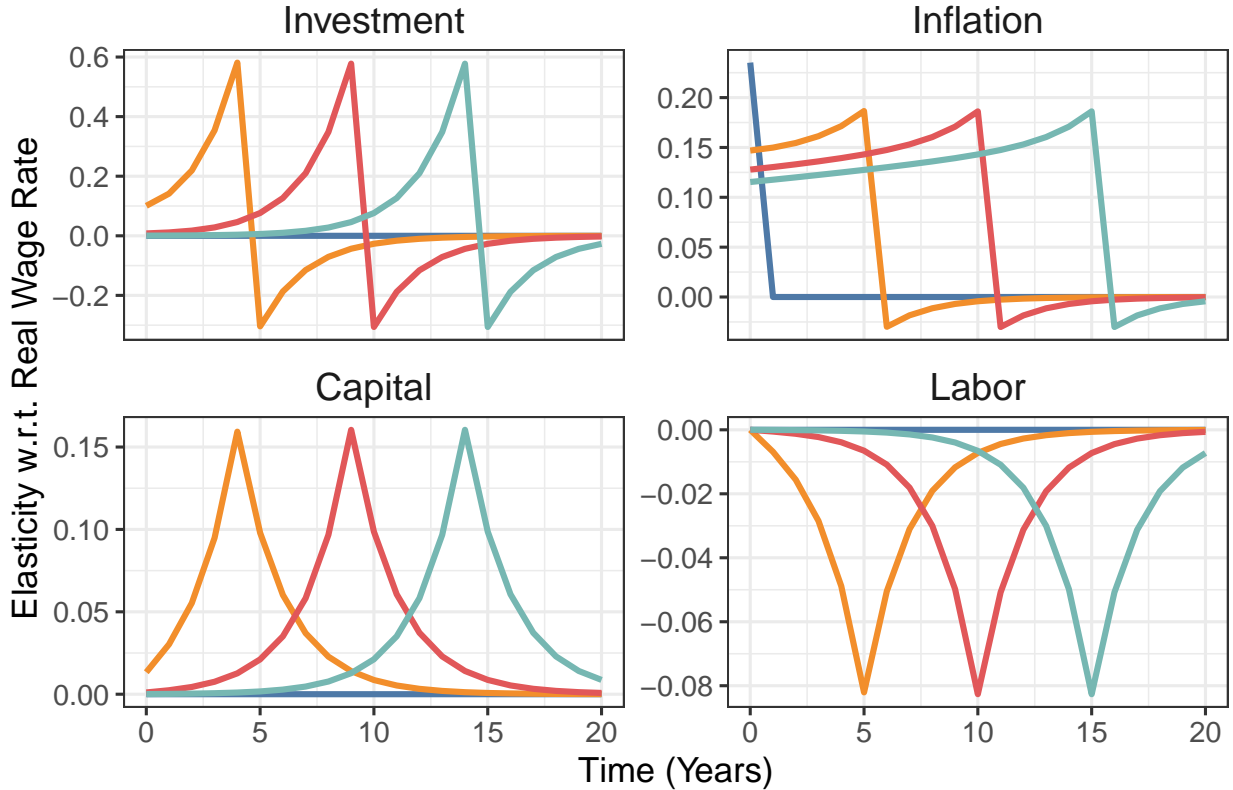


Figure B.3: Jacobians of the Firm Block with Respect to Real Wage

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

$$L_{kt} = \left(\frac{W_{kt}}{W_t} \right)^{-\epsilon_w} L_t. \quad (53)$$

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

$$W_t = \left(\int_0^1 W_{kt}^{1 - \epsilon_w} dk \right)^{\frac{1}{1 - \epsilon_w}}. \quad (54)$$

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

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

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

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

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

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

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

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

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

$$(\pi_t^w - 1) \pi_t^w = \frac{w_t}{\chi_w} \int \dots d\hat{D}_t + \beta (\pi_{t+1}^w - 1) \pi_{t+1}^w. \quad (55)$$

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

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

Thus, the partial is

$$\frac{\partial l_{i,t}}{\partial w_{k,t}} = \frac{\partial \gamma(\mathbf{a}_{i,t}, L_{k,t})}{\partial L_{k,t}} \frac{\partial L_{k,t}}{\partial w_{k,t}} = \frac{\partial \gamma(\mathbf{a}_{i,t}, L_{k,t})}{\partial L_{k,t}} \left(-\epsilon_w \frac{L_{k,t}}{w_{k,t}} \right), \quad (56)$$

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:

$$z_{i,t}(w_{k,t}) \equiv T(\tilde{y}_{i,t} \times w_{k,t} \times l_{i,t}), \quad (57)$$

which is

$$\frac{\partial z_{i,t}}{\partial w_{k,t}} = T'(y_{i,t})\tilde{y}_{i,t} \left[l_{i,t} + w_{k,t} \frac{\partial l_{i,t}}{\partial w_{k,t}} \right] \quad (58)$$

$$\frac{\partial z_{i,t}}{\partial w_{k,t}/w_t} = T'(y_{i,t})y_{i,t} \left[1 + \frac{\partial l_{i,t}/l_{i,t}}{\partial w_{k,t}/w_t} \right] = T'(y_{i,t})y_{i,t} \left[1 - \epsilon_w \frac{L_t}{l_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} \right] \quad (59)$$

Plug these back into (55) to get

$$\begin{aligned} (\pi_t^w - 1) \pi_t^w &= \frac{1}{\chi_w} \left[\int L_t v'(l_{i,t}) \epsilon_w \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} d\hat{D}_t \right. \\ &\quad \left. - \int u'(c_{it}) T'(y_{i,t}) y_{i,t} \left(\epsilon_w \frac{L_t}{l_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} - 1 \right) d\hat{D}_t \right] \\ &\quad + \beta (\pi_{t+1}^w - 1) \pi_{t+1}^w. \end{aligned} \quad (60)$$

Scaling the Phillips curve. Let's define distributional aggregates

$$u'(C_t^*) = \int u'(c_{it}) T'(y_{i,t}) y_{i,t} \left(\epsilon_w \frac{L_t}{l_{i,t}} \frac{\partial \gamma(\mathbf{a}_{i,t}, \tilde{L}_t)}{\partial L_t} - 1 \right) d\hat{D}_t, \quad (61)$$

$$v'(L_t^*) = \int L_t v'(l_{i,t}) \epsilon_w \frac{\partial \gamma(\mathbf{a}_{i,t}, L_t)}{\partial L_t} d\hat{D}_t. \quad (62)$$

Then, the nonlinear Phillips curve becomes

$$(\pi_t^w - 1) \pi_t^w = \frac{1}{\chi_w} \left[v'(L_t^*) - u'(C_t^*) \right] + \beta (\pi_{t+1}^w - 1) \pi_{t+1}^w. \quad (63)$$

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 (63). Let hatted variables denote log-deviations from steady state, assuming that gross inflation in steady state is 1, we get

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

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

$$\hat{\pi}_t^w = \frac{v'(L^*)}{\chi_w} \left[\frac{v''(L^*) L^*}{v'(L^*)} \hat{L}_t^* - \frac{u''(C^*) C^*}{u'(C^*)} \hat{C}_t^* \right] + \beta \hat{\pi}_{t+1}^w, \quad (64)$$

$$\hat{\pi}_t^w = \underbrace{\frac{v'(L^*)}{\chi_w}}_{\kappa_w} \left[\hat{L}_t^* + \rho \hat{C}_t^* \right] + \beta \hat{\pi}_{t+1}^w, \quad (65)$$

where $\nu > 0$ is the reciprocal of the Frisch elasticity, and $\rho > 0$ is relative risk aversion. Equation (65) 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

$$\kappa_w = \frac{1}{1 + \Gamma_w} \frac{[1 - \beta(1 - \xi_w)]\xi_w}{1 - \xi_w}, \quad (66)$$

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 (63) as

$$(\pi_t^w - 1) \pi_t^w = \kappa_w \left(\frac{v'(L_t^*)}{u'(C_t^*)} - 1 \right) + \beta (\pi_{t+1}^w - 1) \pi_{t+1}^w \quad (67)$$

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

Finally, we note that

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

B.3 DAG Representation

In this section, we present the directed acyclic graph (DAG) representation of the full macro model. See Auclert et al. (2021) for a formal introduction of the DAG concept.

- Unknowns: $\{R_t^b, Y_t, \lambda_t, w_t\}$. Exogenous: $\{\varepsilon_t^{mp}\}$.
- Compute the ex-ante rate $\{R_t^b\} \rightarrow \{R_t^e\}$ as $R_t^e = \mathbb{E}_t[R_{t+1}^b]$.
- Evaluate production block as described in Section B.1.3

$$\{Y_t, w_t, R_t^e\} \rightarrow \{L_t, K_t, I_t, Q_t, mc_t, \pi_t, p_t, d_t, R_t^s\}.$$

Note that $p_t^I \equiv 1$ because investment good is the same as the final good in this model.

- Evaluate monetary block $\{\pi_t, \varepsilon_t^{mp}\} \rightarrow \{R_t^n\}$.
- Evaluate household block $\{w_t, L_t, R_t^b, R_t^s, \lambda_t\} \rightarrow \{C_t, A_t^s, A_t^b, T_t, \Lambda_t, v'(L_t^*), u'(C_t^*)\}$.
- Evaluate the labor block $\{v'(L_t^*), u'(C_t^*)\} \rightarrow \{\pi_t^w\}$.

- Solve fiscal / financial intermediary block $\{A_t^s, A_t^b, p_t, R_t^s, R_t^b, T_t, \Lambda_t\} \rightarrow \{B_t\}$ as follows.

– Unknown: $\{B_t\}$.

– Evaluate intermediary block to get $\{d_t^{FI}, N_t\}$ as

$$d_t^{FI} = \iota d^{FI} + (1 - \iota) [R_t^s p_{t-1} + R_t^b B_{t-1} - R_t^s A_{t-1}^s - R_t^b A_{t-1}^b - N], \quad (68)$$

$$N_t = R_t^s p_{t-1} + R_t^b B_{t-1} - R_t^s A_{t-1}^s - R_t^b A_{t-1}^b - d_t^{FI}. \quad (69)$$

– Evaluate fiscal block to get $\{S_t\}$ as

$$S_t = T_t - G_t + \Lambda_t - \mathcal{E} + d_t^{FI}. \quad (70)$$

– Target:

$$0 = B_t + S_t - R_t^b B_{t-1}. \quad (71)$$

- Targets:

$$0 = \frac{R_{t-1}^n}{\pi_t} - R_t^b, \quad (72)$$

$$0 = \pi_t^w \frac{w_t}{w_{t-1}} - \pi_t, \quad (73)$$

$$0 = \lambda_t - \lambda_{ss} + \phi \frac{B_{t-1} - B_{ss}}{Y_{ss}}, \quad (74)$$

$$0 = p_t + B_t - N_t - A_t^s - A_t^b. \quad (75)$$

- Validate Walras's law on (29).

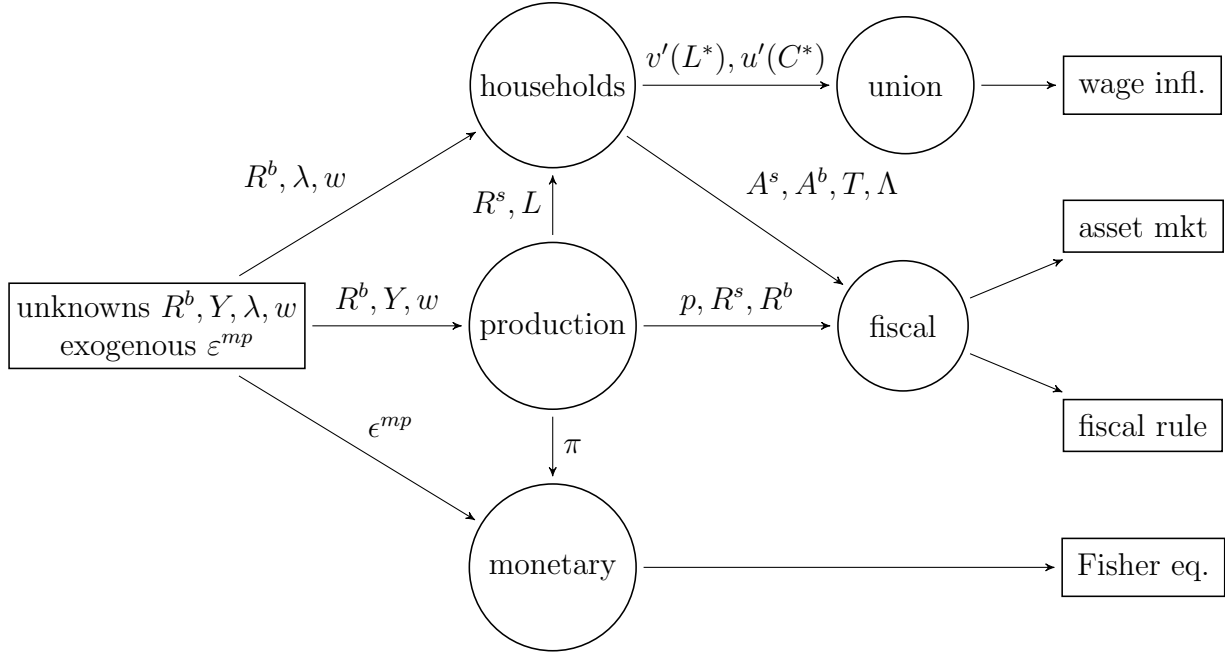


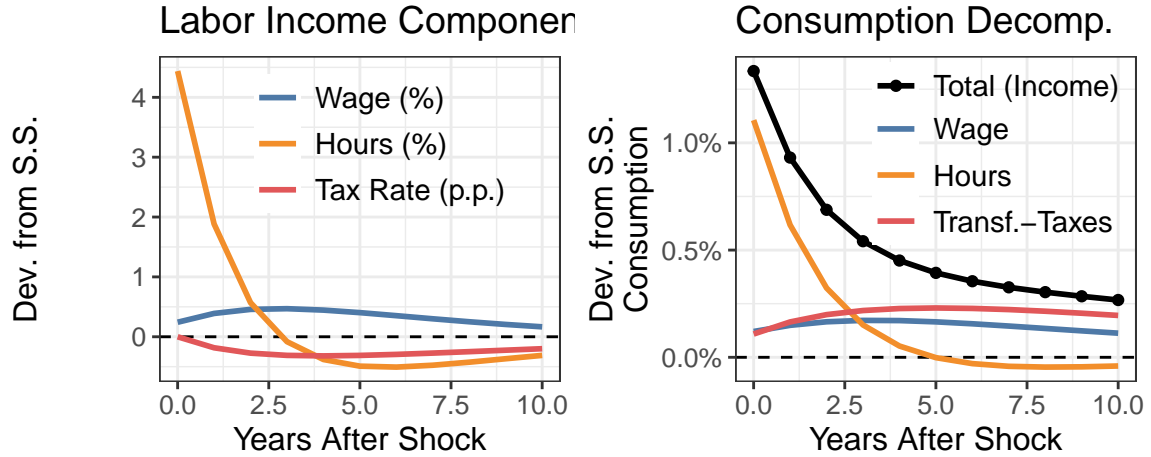
Figure B.4: Directed Acyclic Graph Representation of the HANK model

C Additional Tables and Figures

Table C.1: Estimated Household Parameters

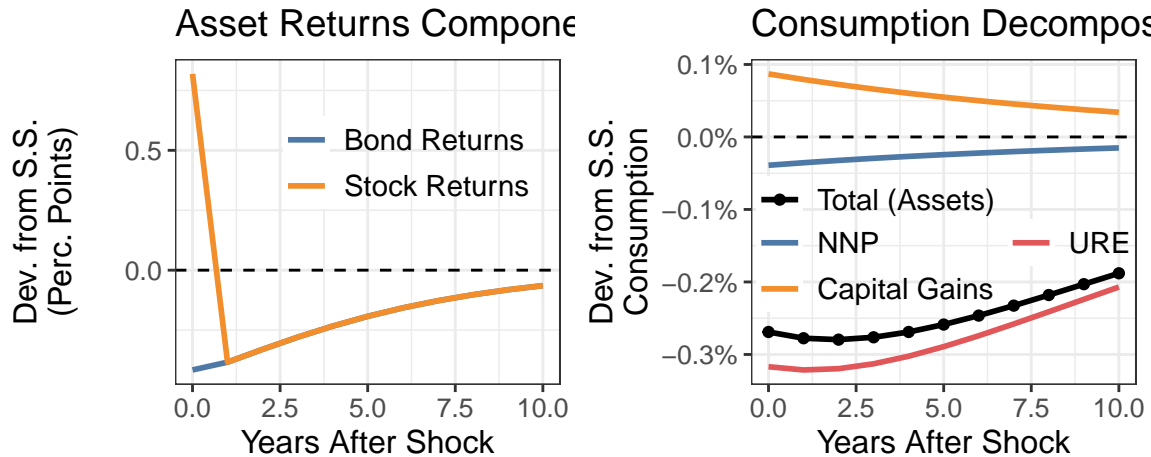
Parameter	Symbol	Estimate
Preferences		
Relative-Risk Aversion	ρ	1.88
Discount Factor	β	0.92
Bequest Intensity	b	299.99
Bequest Shifter	κ	11.00
Income Process		
Income F.E. Std. Dev.	σ_α	0.62

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.



“Wage” denotes the real wage w . “Hours” denotes aggregate labor hours L . “Tax Rate” is $1 - \lambda$ from Equation 5. 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 C.1: Decomposition of Aggregate Consumption Response to a Monetary Policy Shock: Labor Income



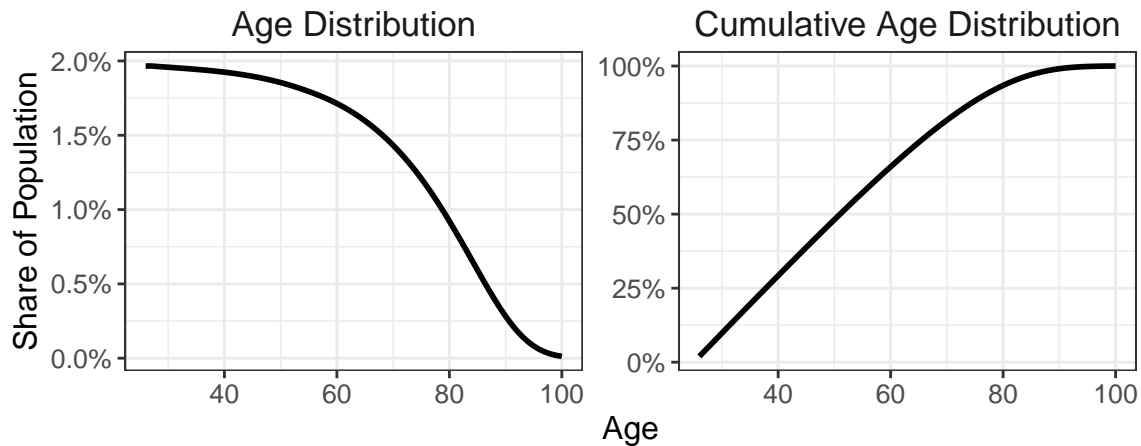
The left panel presents the trajectory of the realized returns of bonds R^b and stocks R^s . They differ only in period 0, when the unanticipated monetary policy shock hits. The right panel presents the response of aggregate consumption when different components of the asset returns channel are given 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 C.2: Decomposition of Aggregate Consumption Response to a Monetary Policy Shock: Asset Returns

Table C.2: 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	1.76
Capital Ratio	K/Y	2.23
Investment Ratio	I/Y	0.30
Dividends Ratio	d/Y	0.12
Government Spending Ratio	G/Y	0.23
Government Debt Ratio	B/Y	0.46
Frisch Elasticity	$1/\nu$	0.50
Parameters		
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	0.91
Capital Dep. Rate	δ_k	0.13

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



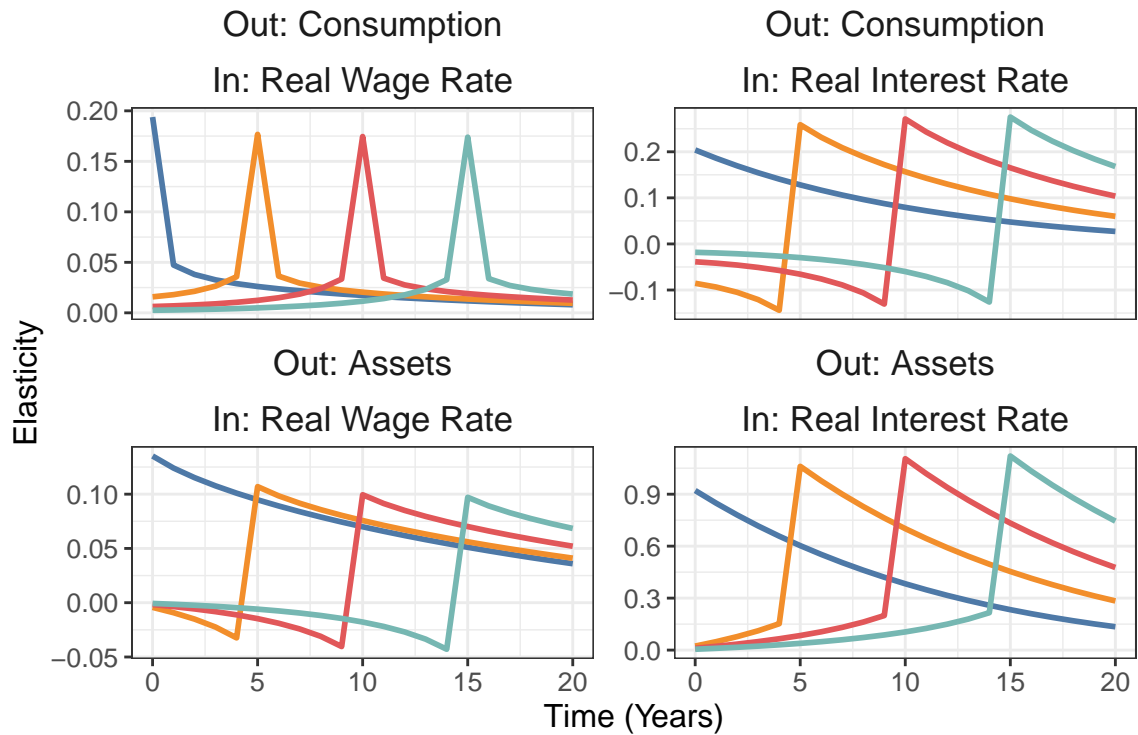
This figure depicts the age distribution of our simulated populations. These come from sequentially applying survival probabilities in the SSA life tables to an initial mass of agents at age 26. The left panel displays, for each age, the share of the population that has that age at any time. The right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

Figure C.3: Age Distribution of Households




See Section 2 for the definitions of each model specification.

Figure C.4: 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.5: 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 32). 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.6: Welfare Effects of a Monetary Policy Shock Net of Disutility of Labor Supply