Portfolio composition and monetary policy transmission

Alicia Zhang
Princeton*

November 1, 2026

[Click here for the latest version]

Abstract

This paper studies the transmission of monetary policy across households that hold heterogeneous portfolios. I derive a sufficient statistics decomposition of the consumption effect of monetary policy that accounts for household-level holdings of assets and liabilities that are differentially exposed to monetary shocks. In this framework, the joint distribution of household portfolio exposures, marginal propensities to consume (MPCs), and changes in asset prices in response to the shock determines the magnitude and distribution of monetary policy transmission. I estimate this decomposition using U.S. data on household balance sheets, MPCs, and asset prices. I find that equity revaluation is a key transmission channel as equity prices respond strongly to monetary shocks. Because equity ownership is concentrated among richer households, expansionary monetary policy generates substantial wealth gains at the top of the wealth distribution. These gains lead high-wealth households to account for a large share of the aggregate consumption response despite their lower MPCs. To assess whether quantitative models capture this distributional pattern, I compare a benchmark two-asset heterogeneous agent New Keynesian model with a parsimonious extension calibrated to moments from the wealth distribution and the equity price response to monetary shocks. The extended model more closely matches the concentration of wealth gains and consumption responses at the top of the wealth distribution.

*Alicia Zhang: aliciazhang@princeton.edu. I am grateful to John Grigsby, Richard Rogerson, Karthik Sastry, and Gianluca Violante for their guidance and support. I thank Mark Aguiar, Moritz Lenel, Atif Mian, Aleksei Oskolkov, and seminar participants at the San Francisco Fed, the New York Fed, and the student macro and macrofinance workshops at Princeton for their comments.

1 Introduction

Monetary policy has swift and striking effects on asset markets. Because households hold markedly different portfolios, the same monetary policy shock can generate very different gains and losses across households. These differences have implications for both the distributional and aggregate effects of monetary policy. This paper takes household portfolio heterogeneity seriously and asks: How does portfolio composition affect individual and aggregate responses to monetary policy?

To answer this question, I derive a set of sufficient statistics for the consumption response to monetary policy shocks. Because asset prices respond differentially to monetary shocks, the aggregate consumption response is determined by the joint distribution of household marginal propensities to consume (MPCs) and changes in portfolio values. Estimating the transmission channels on U.S. data, I find that the largest channel of monetary transmission is the *revaluation channel*, or the change in the value of household portfolios immediately following the shock. This channel is driven by a substantial change of equity prices in response to the shock. However, wealthy households hold the majority of equity in the economy and have lower marginal propensities to consume out of portfolio revaluations, *dampening* the total effect of monetary transmission.

The role of portfolio composition can be illustrated by comparing two households who are identical except for their portfolio allocations. Household A holds all of their wealth in equity, while household B has only bank deposits. Monetary policy shocks tend to have large effects on stock prices (e.g. Bernanke and Kuttner (2005)) and small effects on deposits interest rates (e.g. Kang-Landsberg and Plosser (2022)). A monetary shock would thus affect the value and future returns of Household A's portfolio more strongly than that of Household B. The resulting change in aggregate consumption would further depend on how likely each household is to change their spending in response to wealth shocks.

The paper proceeds in three steps. First, I present a simple theory of household consumption with heterogeneous portfolios and derive a first-order decomposition of the consumption response to a monetary policy shock. While standard mechanisms such as the income and Fisher channels are present in this decomposition, this derivation highlights the importance of households' portfolio compositions in determining their exposure to the asset price and dividend effects of monetary policy shocks. This decomposition separates asset transmission channels dependent on a single reference rate as studied in previous work from *portfolio composition* channels driven by asset markets that respond heterogeneously to the monetary shock. In particular, the joint distribution over household MPCs and changes in portfolio values is a key object that emerges when allowing for portfolio heterogeneity.

Second, I estimate the sufficient statistics from the theory using U.S. data. I discipline household portfolios using detailed data from the Survey of Consumer Finances. For each asset class held by households in the SCF, I run impulse responses of these asset classes to identified monetary shocks and merge the estimates back onto household portfolios to arrive at household-specific portfolio

responses to a monetary shock. This provides an estimate of the household's portfolio exposure to monetary policy. Then, I estimate income, consumption, and MPCs following the procedure in Patterson (2023), and map these back to SCF households by matching on observables. Combining these measures of household portfolio exposures and MPCs provides estimates of each transmission channel of monetary policy.

I find that taking into account realistic household portfolio compositions *amplifies* the transmission of monetary policy, and that portfolio *revaluation* comprises a majority of the consumption response to a monetary policy shock. This effect reflects two opposing forces. First, taking into account empirically observed household portfolios results in a larger estimation of the consumption response compared with a naïve approach of only incorporating the portfolio of a representative household. Second, the portfolio sorting effect is *negative*, as households that are more exposed to the shock have lower MPCs. This is because high-wealth households place disproportionately more money in equities, which are much more price-sensitive to monetary policy shocks than most other assets. Because high-wealth households are so much more exposed to the shock and already account for a large share of aggregate consumption, the aggregate transmission of monetary policy shocks largely flows through portfolio revaluations for low-MPC, high-wealth households. This implies that the portfolio-driven effects of monetary policy are distributed unequally across households.

Finally, I validate this decomposition and compare the size and distribution of consumption responses in two different versions of a two-asset HANK model. The first is a canonical two-asset HANK model following Auclert et al. (2021). In response to empirically relevant monetary policy shocks, the empirical sufficient statistics procedure closely approximates the nonlinear household response. However, this canonical model does not match key moments of the distributional effects of monetary shocks: Wealth gains and consumption response following a monetary expansion are far more uniformly distributed than in the data, which shows greater wealth and consumption increases at the top of the wealth distribution.

To better match these empirical moments while staying close to the canonical model, I introduce a parsimonious extension that better matches moments from the wealth distribution and equity price response to monetary shocks in the data. I augment the canonical model with permanent wealth types, where I calibrate household patience and equity holdings to match wealth and equity shares in the data. In addition, I add an equity valuation wedge motivated by financial intermediary portfolio constraints as in Gabaix and Koijen (2021), and match the impulse response of equity prices from the model to the data. Although the distribution of the consumption response to a monetary shock is not targeted by these moments, this extended model better matches the concentration of wealth and consumption gains at the top of the wealth distribution.

I then use these models to study how the model extensions affect the aggregate transmission of monetary and fiscal policy. I compare impulse responses of aggregate variables following an expansionary monetary policy shock. Consistent with the empirical findings, matching the equity price response amplifies the consumption and inflation effects of monetary policy. Additionally,

matching the wealth distribution dampens the response as high-wealth, low-MPC households have outsize equity holdings. Together, these results suggest that matching moments of the wealth distribution and equity are crucial to accurately modeling the aggregate and distributional effects of monetary policy.

Literature This paper connects with three main strands of the literature. First, the sufficient statistics strategy in this paper extends that in Auclert (2019), who derives a first-order decomposition of the consumption response to a monetary shock. I extend his approach to accommodate empirically observed paths of prices and dividends at the asset class level, which allows households with heterogeneous portfolios to experience different changes in their portfolio values and returns in response to shocks. I derive a household and aggregate consumption decomposition equation to estimate in the data, which nests the estimation equation in Auclert (2019) under specific assumptions. This extension makes the composition of household wealth central to the aggregate and distributional effects of monetary policy: Transmission depends both on the assets that household hold, on how those assets respond to monetary shocks, and on *which* households hold those assets.

Del Canto et al. (2024) similarly derive and estimate sufficient statistics for household responses to shocks, focusing on inflationary shocks and household budget sets rather than consumption responses. I follow a similar methodology as Del Canto et al. (2024) in estimating portfolio responses to monetary shocks through aggregating the impulse responses of individual asset classes. Estimating consumption responses requires combining portfolio IRFs with household marginal propensities to consume (MPCs) as well as carefully addressing the schedule of asset maturation, both of which are not necessary in Del Canto et al. (2024)'s feasible set approach. Other work using sufficient statistics to study monetary policy includes Slacalek et al. (2020), which analyzes transmission in the euro area, and Doepke and Schneider (2006), which studies the winners and losers from inflation. Compared to Slacalek et al. (2020), I derive a unified decomposition equation and an empirical strategy that estimates sufficient statistics at the household level. This allows me to flexibly aggregate households into different groups based on observables, and robustly study the effects of different assumptions in the implementation.

Second, this paper contributes to the quantitative literature that uses HANK models to study monetary policy transmission. Kaplan et al. (2018) pioneered the analysis of monetary policy through HANK models with idiosyncratic income risk, and found that the indirect effects of monetary policy that operate through the labor market are particularly important. Recent work such as Kekre and Lenel (2022) incorporates portfolio heterogeneity into a HANK model through different household propensities to take risk. Melcangi and Sterk (2024) analyze the effect of stock market participation on monetary policy effectiveness. They emphasize the importance of stock market investment and endogenous participation on monetary transmission. In contrast, I focus on a different angle of asset market wealth: Revaluation and cash-flow effects, rather than portfolio adjustments. I also directly estimate the consumption response through asset market movements in the data. Building upon the observation in Del Canto et al. (2024) that standard two-asset HANK

models generate unrealistically small portfolio channels, I add permanent wealth and equity exposure types and an equity valuation wedge disciplined by the data. Similar to Lee (2024), I find that matching the equity price response to monetary shocks is particularly important in generating realistic portfolio responses to a monetary shock.

Finally, this paper connects to a growing literature on heterogeneous household portfolios and returns. Fagereng et al. (2020) report that individuals in Norway earn very different returns on their wealth, and that those with higher net worth earn systematically higher returns. Bach et al. (2020) similarly find that richer Swedish households earn higher returns on their wealth, largely through holding riskier portfolios. Fernández-Villaverde and Levintal (2024) study how changes in asset returns affect inequality in a quantitative model. Studies using administrative data have found heterogeneous monetary policy transmission across different household dimensions: Holm et al. (2021) trace the transmission of policy through Norwegian household balance sheets, Amberg et al. (2022) document U-shaped income effects of monetary shocks in Sweden, and Andersen et al. (2023) show that the gains from expansionary monetary policy are increasing with income. Relative to this work, my contribution is to provide a unified sufficient statistic accounting framework to decompose the effect of monetary policy through various channels, including the portfolio heterogeneity channel. In addition, my method provides a more data-friendly approach to measuring these channels. Each component of the consumption response is estimated from publicly available US data, which allows the decomposition to be recomputed for different shocks and updated in real time based on household portfolio and asset price data.

Roadmap Section 2 derives sufficient statistics for the individual and aggregate consumption responses to a general policy shock. Section 3 describes the data and empirical methodology, and estimates the sufficient statistics decomposition described in the theory. Section 4 quantitatively validates the sufficient statistics approach and studies the distribution of consumption responses in a two-asset HANK model. Section 5 concludes.

2 Sufficient statistics with heterogeneous portfolios

In this section, I develop a framework to study the transmission of monetary policy when individuals hold heterogeneous portfolios. I build on the sufficient statistics approach of Auclert (2019) to study the individual and aggregate response of consumption to a shock that affects economy-wide prices, dividends, and incomes. Households enter the period of the shock holding predetermined asset shares, and their portfolio values and dividends are affected heterogeneously by a monetary shock that flows through differently to each asset. I split the aggregate consumption response into a *reference* component based on a single economy-wide rate, such as a Treasury yield or the return on the average household portfolio, and an *idiosyncratic* component that captures heterogeneity across household portfolio allocations. The key object in this decomposition is the joint expectation between MPCs and portfolio responses to the shock.

2.1 The household problem with heterogeneous portfolios

Environment, assets, and the household problem Time is discrete and indexed by t. Households are continuous with unit mass and are indexed by $i \in \mathcal{I}$. There are $k = \{1, \ldots, K\}$ assets with nominal prices $\mathbf{Q_t} = (Q_t^1, \ldots, Q_t^K)$ and dividends $\mathbf{D_t} = (D_t^1, \ldots, D_t^K)$. To account for both assets and liabilities, I specify that Q_t^k and D_t^k are always non-negative by convention while household holdings of each asset a_{it}^k are positive if the holding is an asset and negative if the holding is a liability.

Importantly, I take household portfolio compositions as given in the following decomposition. This can be microfounded by participation costs (e.g. Vissing-Jorgensen (2022)), search costs (e.g. Menzio and Spinella (2025)), or financial knowledge (e.g. Lusardi et al. (2017)). Lemma 2 shows the assumptions under which this fixed portfolio benchmark, pinned down by a general portfolio constraint, is equivalent to an economy with endogenous portfolio shares.

Assumption 1 (Fixed household portfolio compositions)**.** *Households can only buy or sell an aggregated asset with fixed quantity proportions ω_i^k in each asset k. Conditional on this portfolio composition, households can freely adjust the amount that they save or borrow at the margin.*

Household i enters time 0 with predetermined asset holdings $\mathbf{a_{i0}}$, allocated across assets exogenously given Assumption 1, and chooses consumption c_{it}, labor n_{it}, and future asset positions to maximize discounted utility under perfect foresight:

$$\begin{aligned} \max_{c_{it}, n_{it}, \mathbf{a}_{it+1}} & \sum_{t=0}^{T} \beta_i^t \left[u(c_{it}) - v(n_{it}) \right] \\ s.t. \quad & P_t c_{it} + \mathbf{Q}_t \mathbf{a}_{it+1} = W_{it} n_{it} + y_{it} + (\mathbf{Q}_t + \mathbf{D}_t)\mathbf{a}_{it}, \\ & \mathbf{a}_{it+1} = \boldsymbol{\omega}_i a_{it+1} \end{aligned} \tag{1}$$

where P_t is the price of the consumption good, W_{it} is the idiosyncratic wage, and y_{it} is unearned income. The household takes the paths of prices, wages, and unearned income as given. As in Auclert (2019), I assume that utility from consumption u and disutility from work v are separable, where $u', v' > 0$, $u'' < 0$, and $v'' > 0$. The terminal condition is $\mathbf{a}_{iT+1} = 0$. Appendix A.1 gives the full optimality conditions.

Under Assumption 1, a household's portfolio holdings a_{it} measure units held in the aggregated household-specific asset with fixed quantity proportions ω_i^k, with scalar portfolio prices Q_{it} and dividends D_{it} defined as:

$$\begin{aligned} a_{it}^k = \omega_i^k a_{it}, \qquad & \sum_{k=1}^{K} \omega_i^k = 1, \\ Q_{it} = \sum_{k=1}^{K} \omega_i^k Q_t^k, \qquad & D_{it} = \sum_{k=1}^{K} \omega_i^k D_t^k. \end{aligned} \tag{2}$$

Substituting these portfolio terms into the budget constraint of the household's problem in Equation (1) gives the scalar budget constraint:

$$P_t c_{it} + Q_{it} a_{it+1} = W_{it} n_{it} + y_{it} + (Q_{it} + D_{it}) a_{it}. \tag{3}$$

Note that because Assumption 1 fixes portfolio composition but does not impose additional binding constraints on saving or access to portfolio payouts, resources are fungible at the margin. That is, holding intertemporal prices fixed, the household's consumption depends on income and dividends through their combined present value only.

Lemma 2 gives sufficient conditions under which allowing households to adjust their portfolio composition results in the same first-order consumption response as this fixed composition benchmark.

2.2 Household response to a persistent shock

A stylized monetary shock In the following sections, I consider a shock that generates a sequence of persistent changes to prices and incomes in the economy:

Definition 1 (A stylized monetary shock)**.** *Let ϵ be a monetary shock that realizes at time 0. The effect of ϵ on aggregate equilibrium variables can be summarized by the resulting changes in the paths of prices and incomes:*

$$\{dP_t(\epsilon), \{dY_{it}(\epsilon)\}_{i \in \mathcal{I}}, \{dQ_t^k(\epsilon), dD_t^k(\epsilon)\}_{k \in \mathcal{K}}\}_{t=0}^{T}$$

The following theoretical results hold for any small shock that affects these paths while leaving preferences and feasible asset bundles unchanged. In Section 3, I discipline these paths using impulse responses to identified monetary policy shocks. This experiment generalizes the transitory shock in the main text of Auclert (2019) and is related to the persistent experiment in his Appendix A.3. To derive the sufficient statistics, I take a first-order approximation of household consumption around an optimal baseline allocation.

Resources and intertemporal prices It is useful to define some household budget objects that are important in the consumption decomposition. Many of these terms are analogous to those defined in Auclert (2019) with household-specific portfolio prices and dividends. For notational simplicity, define total real household income as $Y_{it} = (y_{it} + W_{it} n_{it})/P_t$. Next, denote the household-specific real interest rate as

$$R_{it} = \frac{Q_{it} + D_{it}}{Q_{it-1}} \frac{P_{t-1}}{P_t}.$$

Households face idiosyncratic real interest rates due to their portfolio compositions, which lead to differential exposure to asset prices and dividends. The cumulative real return experienced by

household i between time 0 and t is

$$\mathcal{R}_{it} \equiv \prod_{s=1}^{t} R_{is} \quad for \quad t \geq 1, \quad \mathcal{R}_{i0} = 1.$$

A household's portfolio holdings include a dividend payoff as well as a realization value. The *cash flow* component of the household's portfolio is the total real dividend value realized in period t, summed over each asset:

$$CF_{it}^{k} = \frac{D_t^k a_{it}^k}{P_t}, \quad CF_{it} = \sum_k CF_{it}^k.$$

Similarly, the *market value* component of the household's portfolio is the total real market value of assets carried into period t:

$$MV_{it}^{k} = \frac{Q_t^k a_{it}^k}{P_t}, \quad MV_{it} = \sum_k MV_{it}^k$$

All asset prices and dividends in this model are expressed in nominal units, so both are deflated by the price level P_t. A household's total wealth is the sum of their initial resources and discounted disposable income, which also must equal discounted consumption over the time horizon T:

$$\mathcal{W}_{i0} = A_{i0} + \sum_{t=0}^{T} \frac{Y_{it}}{\mathcal{R}_{it}} = \sum_{t=0}^{T} \frac{c_{it}}{\mathcal{R}_{it}}, \quad A_{i0} = MV_{i0} + CF_{i0} \tag{4}$$

The household consumption function Conditional on the path of disposable income including the labor decision n_{it}, a household's consumption allocation problem across time can be written as such due to separability:

$$\max_{\{c_{it}\}} \sum_{t=0}^{T} \beta_i^t u(c_{it}) \quad s.t. \quad \sum_{t=0}^{T} \frac{c_{it}}{\mathcal{R}_{it}} = \mathcal{W}_{i0}$$

Thus, the time-0 consumption function of the household can be written as a function of wealth and intertemporal prices:

$$c_{i0} = \mathcal{C}_i\big(\mathcal{W}_{i0}, \{\mathcal{R}_{it}\}_{t=1}^{T}\big)$$

Given this consumption function, I can derive the household's first-order consumption response to the monetary shock. Under Assumption 1, household i's first-order consumption response at time

0 to the monetary shock specified in Definition 1 is:

$$dc_{i0} = \underbrace{MPC_i d\Omega_{i0}}_{wealth} - \underbrace{c_{i0} \sum_{t=1}^{T} \varepsilon^h_{i0t} d \log \mathcal{R}_{it}}_{substitution} \tag{5}$$

The *wealth effect* depends on the household's marginal propensity to consume MPC_i and their change in cost-adjusted total wealth $d\Omega_{i0}$, defined as:

$$MPC_i = \frac{\partial \mathcal{C}_i}{\partial \mathcal{W}_{i0}},$$
$$d\Omega_{i0} = dA_{i0} + \sum_{t=0}^{T} \frac{dY_{it}}{\mathcal{R}_{it}} + \sum_{t=1}^{T} \frac{c_{it} - Y_{it}}{\mathcal{R}_{it}} d \log \mathcal{R}_{it}$$

Define the intertemporal elasticity of substitution (IES) as $\sigma_{it} = -u'(c_{it})/[c_{it}u''(c_{it})]$. The MPCs out of current income, market wealth, and dividends are *equivalent* in this model, because all resources are perfectly fungible. That is, I can also write the MPC as

$$MPC_i = MPC_i^Y = \frac{\partial c_{i0}}{\partial Y_{i0}} = MPC_i^W \equiv \frac{\partial c_{i0}}{\partial MV_{i0}} = MPC_i^D \equiv \frac{\partial c_{i0}}{\partial CF_{i0}} = MPC_i^{\mathcal{R}} = \frac{\sigma_{i0} c_{i0}}{\sum_{t=0}^{T} \sigma_{it} c_{it} / \mathcal{R}_{it}}. \tag{6}$$

Each MPC is a partial derivative, measuring the change in consumption resulting from the change in one dollar of wealth while holding all other components fixed. Appendix A.4 shows the derivation.

The *substitution effect* depends on the *Hicksian elasticity* ε^h_{i0t} and the *change in intertemporal prices* $d \log \mathcal{R}_{it}$. The Hicksian intertemporal substitution elasticity is defined as:

$$\varepsilon^h_{i0t} = -\frac{\partial \log c^h_{i0}}{\partial \log \mathcal{R}_{it}}$$

where c^h_{i0} is Hicksian consumption. Appendix A.4 derives this decomposition, and Lemma 3 and Assumption 6 further specify the form of ε^h_{i0t}.

Proposition 1 (Household consumption response to a monetary shock)**.** *Household i's first-order consumption response at time 0 to the monetary shock in Definition 1 is:*

$$dc_{i0} = MPC_i \left[\sum_{t=0}^{T} \frac{dY_{it}}{\mathcal{R}_{it}} + dA_{i0} \right] + \sum_{t=1}^{T} \left[MPC_i \frac{c_{it} - Y_{it}}{\mathcal{R}_{it}} - c_{i0} \varepsilon^h_{i0t} \right] d \log \mathcal{R}_{it} \tag{7}$$
$$where \quad dA_{i0} = \frac{1}{P_0} \sum_k a^k_{i0} (dQ^k_0 + dD^k_0) - A_{i0} d \log P_0$$

The proof is in Appendix A.4. Equation 7 contains four main forces. The first term MPC_i $\sum_{t=0}^{T} \frac{dY_{it}}{\mathcal{R}_{it}}$ is the *income channel*, or the change in the household's income path multiplied by their

marginal propensity to consume. The second term $MPC_i dA_{i0}$ is the *portfolio revaluation channel*, or the change in the valuation of the household's portfolio multiplied by their MPC. This term depends on the household's idiosyncratic portfolio composition, the value of their portfolio going into time 0, and the change in prices and dividends for assets held by the household.

The third term $\sum_{t=1}^{T} MPC_i \frac{c_{it}-Y_{it}}{\mathcal{R}_{it}} d\log \mathcal{R}_{it}$ is the *net spending rate exposure*, or the exposure of spending net of income $(c_{it} - Y_{it})$ to changes in the path of real rates $\mathcal{R}_{it}$. The combination of this term and the revaluation term is analogous to the real *unhedged rate exposure* in Auclert (2019). Finally, $-\sum_{t=1}^{T} c_{i0}\varepsilon^{h}_{i0t} d\log \mathcal{R}_{it}$ is the *substitution effect*. This captures how the timing of consumption changes relative to the path of real rates.

Proposition 1 does *not* allow for households to rebalance their portfolios. This follows from Assumption 1: Household portfolio shares are fixed over the course of the shock. This assumption is useful for the empirical implementation, as I cannot observe portfolio adjustments in the SCF, which I use to discipline portfolios. Additionally, this assumption gives a parsimonious, scalar consumption-saving problem while maintaining differences in financial exposure across households.

The following lemmas describe how well the decomposition in Equation 7 holds under endogenous portfolio choice.

Lemma 1 (The revaluation channel is exact irrespective of portfolio choice)**.** *For any predetermined asset holdings $\{a^k_{i0}\}_k$, irrespective of portfolio choice decisions, the change in real financial resources resulting from a change in prices and payoffs is:*

$$dA_{i0} = \frac{1}{P_0}\sum_k a^k_{i0}\left(dQ^k_0 + dD^k_0\right) - A_{i0} d\log P_0$$

Proof. Take the derivative of the household's initial resources defined in Equation (4). □

Lemma 1 establishes that the time-0 revaluation effect, dA_{i0}, does not depend on subsequent portfolio choices. Intuitively, this is because the revaluation effect of a shock is dependent only on a household's portfolio *going into* the shock period, and does not rely on subsequent household decisions. However, the equivalence of the entire consumption decomposition in Equation (7) under endogenous portfolio composition and fixed portfolio shares for other channels relies on more stringent assumptions on feasible portfolio choices.

I explain these assumptions in more detail in Lemma 2. In this Lemma, I consider an extended model where $\mathbf{G}_{it}(\mathbf{a}_{it+1}, \mathbf{a}_{it}) \leq 0$ summarizes general portfolio constraints that are faced by household i.

Lemma 2 (First-order equivalence under endogenous portfolio composition)**.** *Assume that the household's portfolio constraints $\mathbf{G}_{it}(\mathbf{a}_{it+1}, \mathbf{a}_{it}) \leq 0$ are independent of wealth, expenditures, and income. Additionally, assume that a common portfolio composition with weights $\boldsymbol{\omega}_i$ maximizes the return at every date. Then, allowing the tradable portfolio's composition to endogenously adjust instead of being fixed at $\boldsymbol{\omega}_i$ leaves baseline consumption and the first-order response of consumption*

to the ϵ shock unchanged:

$$c_{i0}^{\text{endog}}(\epsilon) - c_{i0}^{\text{fixed}}(\epsilon) = o(|\epsilon|).$$

The proof is in Appendix A.5.The intuition is that if the unconstrained portfolio shares ω_i are already optimal *before* the shock, then an envelope condition mandates that no optimal portfolio adjustments after the shock will change the household's future return paths. This leads consumption to be equivalent to the economy where portfolio constraints imposed the same weights ω_i.

2.3 Aggregation and portfolio composition

Next, I aggregate the household-level responses to arrive at the economy-wide response to a monetary shock. In order to aggregate across heterogeneous portfolio mixes, I define an economy-wide asset as the *reference asset*: The asset that a planner who disregards individual-level portfolio heterogeneity chooses to observe. This makes my results comparable to those in canonical HANK papers where there is a single asset available to households. This does not assume that any households actually earn the reference return; instead, the reference asset is simply the benchmark against which household-specific portfolio changes are defined.

In Section 3, I use the average portfolio composition across households, weighted by wealth as the baseline reference asset. This provides the *representative-agent benchmark*: In a representative agent economy, the agent holds all assets in the economy, so their portfolio composition will be the wealth-weighted portfolio composition. Appendix B.5 additionally shows the results using the zero-coupon Treasury curve from Adrian et al. (2013) as the reference asset. The total change in consumption arising from this decomposition measurement is equivalent regardless of the reference asset; however, the choice of the reference asset determines the exact division of the transmission channels into the reference and heterogeneous channels defined below.

Aggregating across household responses in Equation (7) gives the economy-wide consumption response in Proposition 2.

Proposition 2 (Aggregate response to monetary shocks)**.** *Under continuity assumptions in Appendix A.6, the first order percent change in aggregate consumption in response to the monetary*

shock in Definition 1 is:

$$\frac{dC_0}{C_0} = \underbrace{\frac{1}{C_0}\sum_{t=0}^{T}\mathbb{E}_i\left[MPC_i\frac{dY_{it}}{\mathcal{R}_{it}}\right]}_{income} - \underbrace{\frac{\mathbb{E}_i\left[MPC_i A_{i0}\right]}{C_0}\frac{dP_0}{P_0}}_{Fisher} + \underbrace{\frac{\mathbb{E}_i\left[MPC_i\sum_k|MV_{i0}^k|\right]}{C_0}\frac{d\bar{Q}_0+d\bar{D}_0}{\bar{Q}_0}}_{reference:\ revaluation\ +\ dividends}$$
$$+\underbrace{\sum_{t=1}^{T}\frac{1}{C_0}\mathbb{E}_i\left[MPC_i\frac{c_{it}-Y_{it}}{\mathcal{R}_{it}} - c_{i0}\varepsilon_{i0t}^h\right]d\log\bar{\mathcal{R}}_t}_{reference:\ forward\ returns} \tag{8}$$
$$+\underbrace{\frac{1}{C_0}\mathbb{E}_i\left[MPC_i\left(\sum_k|MV_{i0}^k|\right)\left(\zeta_{i0}^Q+\zeta_{i0}^D\right)\right]}_{composition:\ revaluation\ +\ dividends} + \underbrace{\frac{1}{C_0}\sum_{t=1}^{T}\mathbb{E}_i\left[\left(MPC_i\frac{c_{it}-Y_{it}}{\mathcal{R}_{it}} - c_{i0}\varepsilon_{i0t}^h\right)\zeta_{it}^{\mathcal{R}}\right]}_{composition:\ forward\ returns}$$

Where $C_0 \equiv \mathbb{E}_i[c_{i0}]$ is aggregate consumption, and $\zeta_{i0}^X,\ X \in \{Q, D, \mathcal{R}\}$ are changes in the household's portfolio valuation, payoffs, and returns relative to the reference rate due to idiosyncratic portfolio composition*:*

$$\begin{aligned}
\zeta_{i0}^Q &\equiv \frac{\sum_k a_{i0}^k dQ_0^k}{\sum_k Q_0^k|a_{i0}^k|} - \frac{d\bar{Q}_0}{\bar{Q}_0} \\
\zeta_{i0}^D &\equiv \frac{\sum_k a_{i0}^k dD_0^k}{\sum_k Q_0^k|a_{i0}^k|} - \frac{d\bar{D}_0}{\bar{Q}_0} \\
\zeta_{it}^{\mathcal{R}} &\equiv d\log\mathcal{R}_{it} - d\log\bar{\mathcal{R}}_t, \quad t \geq 1
\end{aligned} \tag{9}$$

Appendix A.6 proves this aggregation and its invariance to the reference asset. The first line includes the income, price level, and reference portfolio channels. The third line describes the portfolio channels stemming from individual deviations from the prices, dividends, and returns of the financial asset. I separate out the individual portfolio terms ζ in order to draw a distinction between this decomposition and a naïve decomposition that takes into account an economy-wide reference asset, which would result in the first two lines only of Equation 8.

Lemma 3 (MPCs and Hicksian substitution elasticities)**.** *Under Assumption 1, consider an interior optimum of the household's consumption problem. Define the elasticity of intertemporal substitution σ_{it} as:*

$$\sigma_{it} = -\frac{u'(c_{it})}{c_{it}u''(c_{it})}, \qquad B_i = \sum_{t=0}^{T}\frac{\sigma_{it}c_{it}}{\mathcal{R}_{it}}$$

where B_i is the discounted present value of consumption weighted by the IES. Holding the path of intertemporal prices and income fixed, the marginal propensity to consume out of time 0 wealth

MPC_i and the elasticity of substitution across time t ε^h_{i0t} are:

$$MPC_i = \frac{\sigma_{i0}c_{i0}}{B_i}, \qquad \varepsilon^h_{i0t} = \frac{\sigma_{i0}\sigma_{it}c_{it}}{B_i\mathcal{R}_{it}}, \qquad t = 1, \ldots, T \tag{10}$$

The sum of the substitution elasticity across future dates is:

$$\sum_{t=1}^{T} \varepsilon^h_{i0t} = \frac{\sigma_{i0}}{B_i} \underbrace{\sum_{t=1}^{T} \frac{\sigma_{it}c_{it}}{\mathcal{R}_{it}}}_{B_i - \sigma_{i0}c_{i0}} = \sigma_{i0}\frac{B_i - \sigma_{i0}c_{i0}}{B_i} = \sigma_{i0}(1 - MPC_i) \tag{11}$$

Equation (11) derives the cumulative substitution elasticity: A uniform one percent increase in cumulative returns at all future dates decreases current Hicksian consumption by $\sigma_{i0}(1 - MPC_i)$ percent. Appendix A.3 provides the full proof.

2.4 Special case: Nesting Auclert (2019)

The main decomposition in Equation (8) follows much of the same theoretical methodology as in Auclert (2019), but with a few key differences. First, the shock I consider causes a sequence of deviations from time $t = 0, \ldots, T$, while Auclert (2019) considers a shock at time 0 only in the main text of the paper, with a theoretical extension in his Appendix. Second, I use changes in asset prices from the data to discipline the effect of monetary policy shocks on household portfolios, while the empirical strategy in Auclert (2019) assumes that future asset prices and payoffs are determined completely by changes in the real rate. In my empirical implementation, these differences require measuring the joint expectation between MPCs, household portfolio exposures, and empirical changes in asset returns and dividends following a monetary shock, while the procedure in Auclert (2019) does not require measuring asset return and dividend changes. As shown in Section 3, these asset class changes both amplify the consumption response compared to the Auclert (2019) benchmark and change the distribution of responses across households.

To show that this decomposition nests the main estimation equation from Auclert (2019) under some additional assumptions, start from Equation (8) and the accompanying assumptions. Let $C_0 = \mathbb{E}_i[c_{i0}]$. For this comparison, let Y^g_{i0} be current gross income, so $Y_{i0} = Y^g_{i0} - \tau_{i0}$, and $Y_0 = \mathbb{E}_i[Y^g_{i0}] > 0$. Let $f_{it} = b_{it} + B^n_{it}/P_t$ denote inherited real contractual payments, including maturing principal, excluded from Y_{it}. Here, b_{it} are real assets and B^n_{it} are nominal assets. With common discount factors $q_t = 1/\mathcal{R}_t$, define

$$URE_{i0} = Y^g_{i0} - \tau_{i0} + f_{i0} - c_{i0},$$
$$NNP_{i0} = \sum_{t=0}^{T} q_t B^n_{it}/P_t.$$

Where NNP_{i0} is the net nominal position.

Consider the following *transitory* shock at time 0:

(a) Incomes change by dY^g_{i0}

(b) The price level permanently changes by $\frac{dP_0}{P_0}$

(c) The one-period real return between time 0 and 1 changes by $\frac{dR_1}{R_1}$

In this environment, make the additional assumptions:

(i) Labor supply is fixed at every date and saving is interior.

(ii) Intertemporal prices are common across households both before and after the shock, but may change with the shock.

(iii) The real bonds b_{it} and nominal bonds B^n_{it} are fixed. In a neighborhood of the baseline, inherited resources satisfy $A_{i0} = \sum_{t=0}^{T} q_t f_{it}$.

(iv) Assume the following uniform lump-sum fiscal adjustment:

$$d\tau_{i0} = d\tau_0 = -\mathbb{E}_i[NNP_{i0}]\frac{dP_0}{P_0} + \mathbb{E}_i[URE_{i0}]\frac{dR_1}{R_1}$$

(v) The time 0 IES is common across households, $\sigma_{i0} = \sigma_0$, and the change in income takes the following form

$$dY^g_{i0} = \left[\frac{Y^g_{i0}}{Y_0} + \gamma\left(\frac{Y^g_{i0}}{Y_0} - 1\right)\right] dY_0$$

for a common exposure parameter γ.

Then Equation (8) reduces to Equation (25) in Auclert (2019):

$$\frac{dC_0}{C_0} = (\mathcal{M}_0 + \gamma\mathcal{E}_{Y,0})\frac{dY_0}{Y_0} - \mathcal{E}_{P,0}\frac{dP_0}{P_0} + (\mathcal{E}_{R,0} - \sigma_0\mathcal{S}_0)\frac{dR_1}{R_1} \tag{12}$$

where

$$\mathcal{M}_0 = \frac{\mathbb{E}_i[MPC_{i0}Y^g_{i0}]}{C_0}, \qquad \mathcal{E}_{Y,0} = \frac{\mathrm{Cov}_i(MPC_{i0}, Y^g_{i0})}{C_0},$$
$$\mathcal{E}_{P,0} = \frac{\mathrm{Cov}_i(MPC_{i0}, NNP_{i0})}{C_0}, \qquad \mathcal{E}_{R,0} = \frac{\mathrm{Cov}_i(MPC_{i0}, URE_{i0})}{C_0},$$
$$\mathcal{S}_0 = \frac{\mathbb{E}_i[(1 - MPC_{i0})c_{i0}]}{C_0}$$

Proof. The household's lifetime budget constraint implies

$$\sum_{t=1}^{T} \frac{c_{it} - Y_{it}}{\mathcal{R}_{it}} = A_{i0} + Y_{i0} - c_{i0}.$$

Under the transitory experiment, changes in the household's total wealth at time 0 are:

$$dA_{i0} + dY_{i0} = dY^g_{i0} - d\tau_{i0} - (A_{i0} - f_{i0})\frac{dR_1}{R_1} - NNP_{i0}\frac{dP_0}{P_0}.$$

Differentiating the pricing identity $A_{i0} = \sum_{t=0}^{T} q_t f_{it}$ under the specified shock gives $dA_{i0} = -NNP_{i0}\frac{dP_0}{P_0} - (A_{i0} - f_{i0})\frac{dR_1}{R_1}$. Combining these expressions yields:

$$\begin{aligned} d\Omega_{i0} &= dA_{i0} + dY_{i0} + \sum_{t=1}^{T} \frac{c_{it} - Y_{it}}{\mathcal{R}_{it}}\frac{dR_1}{R_1} \\ &= dY^g_{i0} - d\tau_{i0} - NNP_{i0}\frac{dP_0}{P_0} + (A_{i0} - (A_{i0} - f_{i0}) + Y_{i0} - c_{i0})\frac{dR_1}{R_1} \\ &= dY^g_{i0} - d\tau_{i0} - NNP_{i0}\frac{dP_0}{P_0} + URE_{i0}\frac{dR_1}{R_1}, \end{aligned}$$

where the final equality uses $URE_{i0} = Y_{i0} + f_{i0} - c_{i0}$. The valuation step follows from the pricing identity. As in Lemma 3, the Hicksian elasticities satisfy $\sum_{t=1}^{T} \varepsilon^h_{i0t} = \sigma_0(1 - MPC_{i0})$. The household consumption response is therefore:

$$dc_{i0} = MPC_{i0}\left(dY^g_{i0} - d\tau_{i0} - NNP_{i0}\frac{dP_0}{P_0} + URE_{i0}\frac{dR_1}{R_1}\right) - \sigma_0 c_{i0}(1 - MPC_{i0})\frac{dR_1}{R_1}.$$

The government's fiscal adjustment is:

$$d\tau_{i0} = -\mathbb{E}_i[NNP_{i0}]\frac{dP_0}{P_0} + \mathbb{E}_i[URE_{i0}]\frac{dR_1}{R_1}$$

Substituting this in and taking expectations converts the exposure terms into covariances:

$$\begin{aligned} dC_0 = \mathbb{E}_i[MPC_{i0}\, dY^g_{i0}] &- \mathrm{Cov}_i(MPC_{i0}, NNP_{i0})\frac{dP_0}{P_0} \\ &+ \{\mathrm{Cov}_i(MPC_{i0}, URE_{i0}) - \sigma_0 \mathbb{E}_i[(1 - MPC_{i0})c_{i0}]\}\frac{dR_1}{R_1}. \end{aligned}$$

Finally, income exposure in Assumption (v) implies:

$$\begin{aligned} \frac{\mathbb{E}_i[MPC_{i0}\, dY^g_{i0}]}{C_0} &= \left\{\frac{\mathbb{E}_i[MPC_{i0} Y^g_{i0}]}{C_0} + \gamma\frac{\mathrm{Cov}_i(MPC_{i0}, Y^g_{i0})}{C_0}\right\}\frac{dY_0}{Y_0} \\ &= (\mathcal{M}_0 + \gamma\mathcal{E}_{Y,0})\frac{dY_0}{Y_0} \end{aligned}$$

Dividing the aggregate consumption response by C_0 and using the remaining definition gives the desired decomposition expression:

$$\frac{dC_0}{C_0} = (\mathcal{M}_0 + \gamma\mathcal{E}_{Y,0})\frac{dY_0}{Y_0} - \mathcal{E}_{P,0}\frac{dP_0}{P_0} + (\mathcal{E}_{R,0} - \sigma_0\mathcal{S}_0)\frac{dR_1}{R_1}.$$

□

3 Monetary transmission channels in U.S. data

Next, I estimate the sufficient statistics from Equation 8 in the data. One key difference between my empirical approach and that in Auclert (2019) is that my moments depend on the *responses* of incomes, prices, and asset prices to identified monetary policy shocks. My empirical strategy proceeds in three steps. First, I construct household portfolios using detailed portfolio composition data from the Survey of Consumer Finances (SCF). Next, I estimate marginal propensities to consume using the Panel Study of Income Dynamics (PSID) and Consumer Expenditure Survey (CEX) following Patterson (2023), and map these MPCs back onto SCF households using a set of observables. Finally, I estimate impulse response functions (IRFs) of incomes, prices, and asset prices to an identified monetary policy shock, and aggregate these IRFs into portfolios for the SCF households. This joint estimation of portfolio exposures, MPCs, and impulse responses provides an accounting of each transmission channel from monetary policy to aggregate consumption.

3.1 Data

My primary sources of household microdata come from the SCF, the PSID, and the CEX. All dollar values are converted to 2010 dollars. I use twelve waves of the SCF from 1989 to 2025 to measure the asset share of broad asset classes that make up each household's portfolio. I keep households whose heads are aged 25-95, with a total of 55,976 households over the entire sample period. The SCF is a particularly good source of data for studying household balance sheets, as the survey records information about a wide variety of household portfolio holdings. Additionally, its intentional oversampling of wealthy households provides better coverage of households that are likely to both have high baseline levels of consumption and be more exposed to financial markets (Bricker et al., 2017).

Additionally, I use data from the PSID and the CEX to estimate marginal propensities to consume following the approach of Patterson (2023). I include households from the PSID in the 1986, 1996, and biennial waves from 2003 to 2021, which leads to a sample size of 9,736 households and 1,227 unemployment events used to instrument for income shocks. Total consumption is imputed from the CEX through a similar method as in Blundell et al. (2008). The CEX sample covers 1981–2023, excluding 1982–1983, and contains 293,777 interviews from 197,481 consumer units. I measure total consumption within the CEX and map predicted consumption back onto the PSID using a variety of household observables.

3.2 Portfolio composition

I construct household portfolio shares in the SCF following a similar procedure as in Fagereng et al. (2020). Let w_i^k be household i's wealth held in category k, where $\mathcal{A}^+$ denotes the set of asset categories and $\mathcal{A}^-$ denotes the set of liability categories. I record liabilities with a negative sign in

wealth w_i^k while assets are positive. A household's share s_i^k in asset k is:

$$s_i^k = \frac{w_i^k}{\sum_{j \in \mathcal{A}^+} w_i^j} \tag{13}$$

As in Fagereng et al. (2020), I define asset shares as a fraction of *gross wealth*, or the sum of assets before subtracting liabilities. This prevents asset shares from becoming extremely large for households with net wealth, which subtracts out liabilities, close to zero. Households with gross wealth of exactly zero are kept in the computation and are assigned income, consumption, and MPC estimates, but their asset-share ratios are undefined. Zero gross assets do not rule out debt, and I maintain liability-related transmission channels for these households using level exposures. These households account for about 2 percent of the total sample. Figure 1 shows the average portfolio composition for each net worth percentile in the SCF across time. Appendix B.4 provides details on which asset categories and aggregation procedures are used.

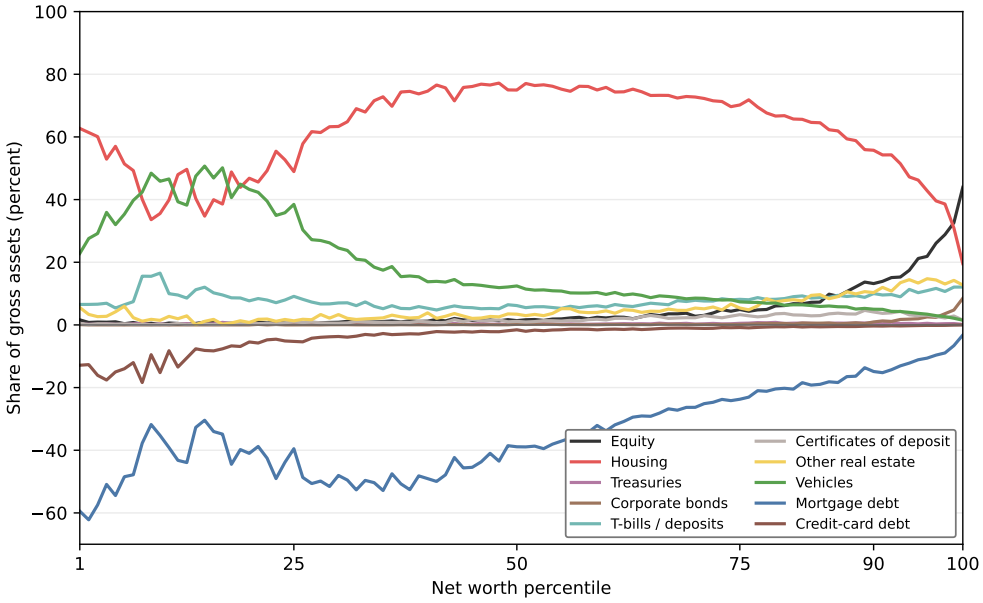


Figure 1: Average portfolio shares in the SCF across the wealth distribution

In general, lower-wealth households hold more debt, middle-wealth households maintain a large portion of their assets in housing, while high-wealth households hold a sizable share of equities. These portfolio composition differences create *household-specific exposure* to monetary policy shocks, as different assets respond differently to the shock.

3.3 Marginal propensities to consume

I measure household-level marginal propensities to consume following Patterson (2023). In this section, I detail some additional assumptions made to map MPCs from the theory to the data,

then describe the estimation procedure.

Resource-specific MPCs The theoretical decomposition in Propositions 1 and 2 assumes that all resources are fungible, and thus households have the same MPC out of all components of wealth. In practice, many papers have estimated different MPCs out of wealth than income (e.g. Di Maggio et al. (2020), Chodorow-Reich et al. (2021)). Let MPC_i, MPC_i^Q, and MPC_i^D be partial derivatives of consumption with respect to changes in income, capital gains, and dividends respectively, holding all else fixed. For the empirical implementation, I allow capital gains and dividends to differ from income in how easily the household can access changes in these accounts after the shock.

Assumption 2 (Spendability constraints on capital gains and dividends)**.** *Following the shock in Definition 1, an additional real dollar of capital gains or dividends increases spendable resources for each household i by a fraction of that dollar, $\lambda_i^Q, \lambda_i^D \in [0,1]$. Income is fully spendable immediately. There are no other changes in the household's objective function or constraints.*

To derive the new MPC on capital gains and dividends, the chain rule gives:

$$MPC_i^Q = \lambda_i^Q MPC_i, \qquad MPC_i^D = \lambda_i^D MPC_i. \tag{14}$$

If resources are fully accessible, then $\lambda_i^Q, \lambda_i^D = 1$ and all MPCs return to the initial MPC out of fully fungible wealth. Note that because income is assumed to be fully spendable, its MPC remains unchanged.

Measuring consumption I follow Patterson (2023) in estimating household marginal propensities to consume. The strategy is to compute the change in consumption associated with a change in earnings caused by an unexpected unemployment shock. While Patterson (2023) estimates this statistic at the individual level, I treat the unit of observation as a household instead to better map household income and consumption to household portfolios in the SCF. I use the sum of earnings from the household head and spouse as the household's income observation and construct the unemployment instrument to be whether either previously employed adult in the household experiences unemployment.

To construct a consistent measure of consumption across the PSID waves, I use the CEX to impute total consumption of each PSID household using information on the relationship between total consumption, food consumption, and household demographics. For each CEX consumer unit j, which I map into PSID households, I regress annual total spending on food spending, a set of household demographic variables, time, and interactions:

$$\log C_{jt} = \alpha + \beta \log F_{jt} + \gamma' (X_{jt}) + \varepsilon_{jt} \tag{15}$$

where C_{jt} is annual total spending, F_{jt} is annual food expenditures, and X_{jt} contains demographic and price index controls, log income, employment indicators, and interactions. I impute total con-

sumption in the PSID following Patterson (2023) by projecting each PSID household i's observables F_{it} and X_{it} on the estimated coefficients and exponentiating:

$$\widehat{C}_{it} = \exp\Big[\widehat{\alpha} + \widehat{\beta}\log F_{it} + \widehat{\gamma}' X_{it}\Big] \tag{16}$$

Appendix B.2 also states an additional IV exclusion restriction for the imputed consumption measure. This gives a measure of annual household consumption for each PSID household.

Estimating MPCs To estimate the MPC of each PSID household, I regress the 2-year change in imputed consumption per household on the corresponding change in earnings and a vector of controls:

$$\Delta C_{it} = (\beta_0 + \boldsymbol{\beta}' X_{it})\Delta E_{it} + \boldsymbol{\alpha}' X_{it} + \gamma \Delta N_{it} + \kappa J_{it-2} + \delta_{t,s(i)} + \varepsilon_{it} \tag{17}$$

where ΔC_{it} is the change in imputed consumption from the household's PSID wave and ΔE_{it} is the corresponding change in earnings. X_{it} is a set of household observables: Liquid wealth terciles, income terciles, household head age, and household head age squared. Liquid wealth and income groups are measured at the previous survey date. ΔN_{it} is any change in the number of people in the household, and J_{it-2} is the count of how many individuals in the household were employed at the last survey date. $\delta_{t,s(i)}$ is a state $s(i)$ $\times$ time t fixed effect for household i. Because the PSID data is collected biennially, all changes are at the two-year horizon. Standard errors are clustered by household. I instrument changes in labor earnings $\{\Delta E_{it}, \Delta E_{it} X_{it}\}$ with unemployment shocks at that period $\{u_{it}, u_{it} X_{it}\}$. The estimated marginal propensity to consume is:

$$\widehat{MPC}_i(X_{it}) = (\hat{\beta}_0 + \hat{\boldsymbol{\beta}}' X_{it}) \tag{18}$$

The identifying assumption is that changes in household employment affect household consumption solely through changes in earnings. Assumption 3 describes the instrument validity and Assumption 4 imposes an additional assumption on the comparability of income shocks across individuals, which is necessary for these shocks to be identifying the same MPCs across households.

Assumption 3 (Unemployment instrument validity and comparable income shocks)**.** *Conditional on household observables and controls, unemployment shocks cause unexpected changes in earnings and affect consumption only through the resulting income changes. A sufficient condition for exogeneity is:*

$$\mathbb{E}[u_{it}\varepsilon_{it} \mid X_{it}, \Delta N_{it}, J_{it-2}, t, s(i)] = 0.$$

Additionally, the instruments $\{u_{it}, u_{it}X_{it}\}$ are jointly relevant for and predict sufficient variation in $\{\Delta E_{it}, \Delta E_{it}X_{it}\}$, to identify the coefficients β_0 and $\boldsymbol{\beta}$.

Assumption 4 (Unemployment shocks induce a common path of earnings losses)**.** *Unemployment shocks u_{it} induce a common time profile of future earnings losses as a percent of the initial shock.*

Specifically,

$$\frac{\mathbb{E}[E_{it+2h}(u_{it}) - E_{it+2h}(0)]}{\mathbb{E}[\Delta E_{it}(u_{it}) - \Delta E_{it}(0)]} = \frac{\mathbb{E}[E_{j,\tau+2h}(u_{j\tau}) - E_{j,\tau+2h}(0)]}{\mathbb{E}[\Delta E_{j\tau}(u_{j\tau}) - \Delta E_{j\tau}(0)]}, \qquad \forall i, j; h = 0, 1, 2, \ldots$$

Assumption 4 fixes the profile of earnings losses caused by an unemployment shock to be constant across households, which implies that the MPC measurement is consistent across households. Figure 2 shows that the persistence of income loss after an unemployment shock is fairly even across liquid wealth by income terciles, which suggests that this assumption holds true in the data. The units of the y-axis are in percentage terms of the income loss following the unemployment shock at time 0, so the period $t = -2$ is normalized to 100% and period $t = 0$ is normalized to 0%. The following periods indicate what percentage of the income loss has been recuperated by each liquid wealth by income group following the unemployment shock. Aside from the high liquid wealth, low income group, which is very noisy due to experiencing only 19 unemployment shocks during this period, the amount of income lost following an unemployment shock is largely homogeneous across all groups. This suggests that the percent change in income, and thus the MPC measured by the unemployment instrument, is comparable across households.

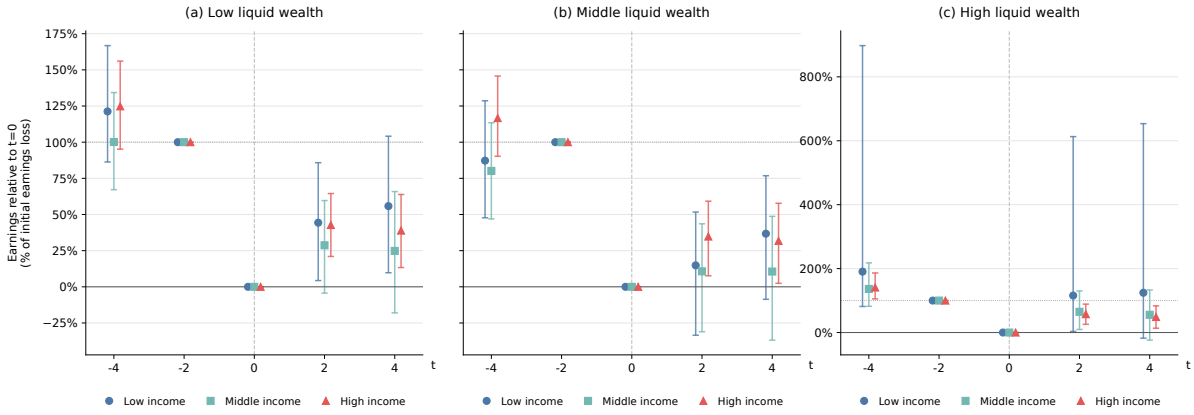


Figure 2: Earnings following an unemployment shock at $t = 0$

Additionally, I compare the unemployment IV estimates with permanent and transitory income responses based on the covariance moments of Blundell et al. (2008). Appendix B.3 reports the comparison on a common PSID sample, including the same conversion from log to dollar consumption responses. The permanent income MPC estimate following the Blundell et al. (2008) methodology is lower on average, but the differences across liquidity and income groups are similar. This suggests that while the level of the MPC measured from the unemployment IV estimate may not exactly match all estimates made using other methodologies, as is commonly found in the literature, the slope in the MPC across groups and thus the distribution of consumption responses due to MPC differences is consistent.

Assigning MPCs to SCF households I take the MPCs estimated in the PSID and map them back to households in the SCF using the same demographic variables in equation (17). The MPC for each SCF household is:

$$MPC_i^{SCF} = \widehat{MPC}(X_{it}^{SCF})$$

Figure 3 shows the average estimated MPC by earnings and liquid net wealth percentiles. In general, MPCs are higher for households with lower labor earnings and lower for higher earners. There is less variation across the liquid wealth distribution, although the average MPC for households with high liquid wealth is slightly lower than households with low liquid wealth.

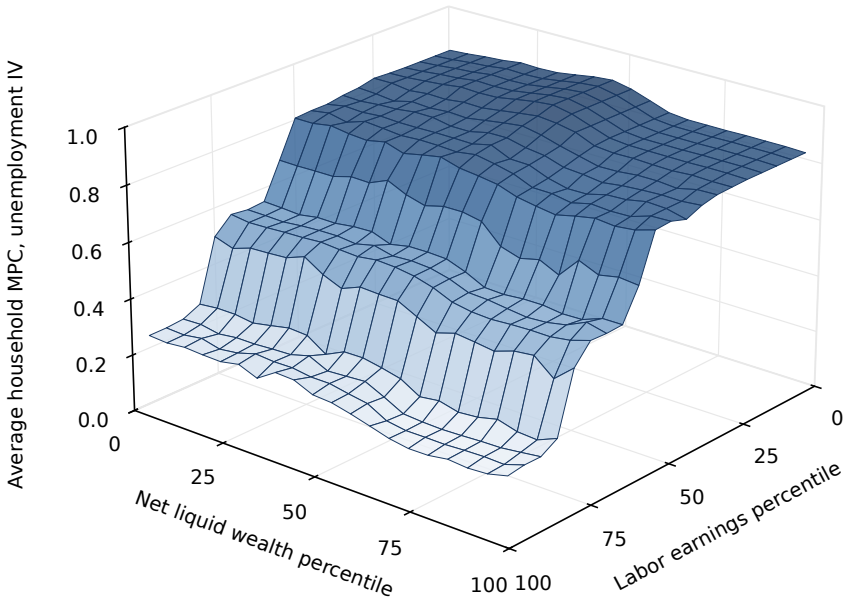


Figure 3: Average household MPCs across the liquid wealth and earnings distribution

Connecting MPCs from income to MPCs from wealth

Assumption 5 (Group-specific spendability shares)**.** *Within each wealth group j, the spendability constraints on capital gains and dividends are equal at λ_j.*

Following Patterson (2023) and Becko et al. (2026), I interpret the coefficient in Equation (18), $\widehat{MPC}_i$, as the household's MPC out of income. Under the assumption that unemployment and monetary policy shocks have comparable time profiles, this MPC is informative about spending responses to income changes generated by monetary policy shocks. Then, under Assumption 5, I use the income, capital-gain, and dividend estimates from Table III in Di Maggio et al. (2020) to calibrate the relative resource-specific MPCs in Equation (14). Let $MPC^s_{DM,j}$ denote the corre-

sponding coefficients for resource s. Then for a household i in group $j(i)$:

$$\lambda_j^Q = \frac{MPC_{DM,j}^Q}{MPC_{DM,j}^Y}, \qquad \lambda_j^D = \frac{MPC_{DM,j}^D}{MPC_{DM,j}^Y}, \qquad MPC_{DM,j}^Y \neq 0,$$
$$\widehat{MPC}_i^Q = \lambda_{j(i)}^Q \widehat{MPC}_i^Y, \qquad \widehat{MPC}_i^D = \lambda_{j(i)}^D \widehat{MPC}_i^Y. \tag{19}$$

Table A.1 documents the exact scaling and resulting MPCs that will be used in the empirical decomposition. This is the empirical extension of the MPCs derived in Proposition 1. The remaining coefficient to be specified is the Hicksian intertemporal elasticity of substitution, which dictates the household's response to future rate changes.

Measuring the Hicksian intertemporal elasticity In addition to requiring an estimate of MPCs, the empirical decomposition also depends on the intertemporal substitution, or how much current consumption responds to cumulative real returns in the future. I make an additional assumption on the elasticity of intertemporal substitution, which will simplify the construction of the Hicksian substitution elasticity.

Assumption 6 (Constant intertemporal substitution)**.** *Household i has a constant elasticity of intertemporal substitution:*

$$\sigma_{it} = \sigma_i > 0, \qquad t = 0, \ldots, T.$$

This elasticity may differ across households.

Combining this assumption on a constant IES with the form of ε_{i0t}^h derived in Lemma 3, the Hicksian elasticity can be written as

$$\varepsilon_{i0t}^h = \sigma_i(1 - MPC_i)\rho_{it}, \qquad \rho_{it} \equiv \frac{c_{it}/\mathcal{R}_{it}}{\sum_{s=1}^T c_{is}/\mathcal{R}_{is}}, \qquad t = 1, \ldots, T. \tag{20}$$

where ρ_{it} weights the contribution of time t to the time 0 substitution response according to its share of discounted future consumption. Because $\sum_{t=1}^T \rho_{it} = 1$, these weights allocate the total substitution elasticity across future dates. I use the estimated income MPC in constructing this measurement of the IES.

3.4 Impulse responses to monetary shocks

For each outcome variable x_t, I estimate impulse responses using the monthly proxy-SVAR approach in Bauer and Swanson (2023). I estimate a separate system containing the outcome, industrial production, CPI, the excess bond premium from Gilchrist and Zakrajšek (2012), and the two-year

Treasury yield. The vector of variables is

$$Y_t(x) = [100 \log IP_t,\ 100 \log CPI_t,\ EBP_t,\ x_t,\ T2Y_t]'.$$

I regress each variable on twelve monthly lags of all five variables, an intercept, and monthly indicators for the COVID period from March–December 2020. Let a denote the vector of intercepts, B_ℓ the matrix of coefficients on the ℓth lag, and d_t the vector of monthly indicators, with coefficient matrix Γ. The system is

$$Y_t(x) = a + \sum_{\ell=1}^{12} B_\ell Y_{t-\ell}(x) + \Gamma d_t + u_t,$$

where u_t is the vector of innovations, estimated as the residuals from these regressions.

I instrument the residual using the monthly monetary surprise series z_t from Bauer and Swanson (2023). This series measures interest rate futures surprises around FOMC announcements and is extended through December 2023. Identification requires that z_t is correlated with monetary policy shocks and uncorrelated with other structural shocks. I estimate the impulse response separately for each asset class from 1983-2023, or a subset of these years if data is not available for the full period.

Figure 4 shows the change in returns for the main set of assets used in estimating portfolio prices and returns for households in the SCF. The shock is normalized to a 25 basis point decrease in the 2-year Treasury yield. Appendix B.5 shows the impulse responses for additional macroeconomic variables. Equities are the most sensitive asset by far to monetary shocks, with a price change over an order of magnitude larger than other assets in the first month of the shock. This is in line with previous work (e.g. Bernanke and Kuttner (2005)) that finds large stock market effects of monetary policy. Expansionary policy also reduces credit card and mortgage interest rates, while decreasing returns for certificates of deposit. Some assets such as housing and vehicle prices do not change in statistically significant ways when using this approach. Since private business wealth makes up a substantial portion of total wealth, I assume that the impulse response of private business wealth is 0.4 $\times$ the impulse response of equity. This comes from the correlation between historical private and public returns from Bricker et al. (2025).[1]

3.5 Sufficient statistics estimates

Combining household portfolios, MPCs, and impulse responses described above, Figure 5 shows the empirical decomposition from Equation (8). Additional details on the implementation are in Appendix B. The baseline estimation is made with the assumption that $\sigma_i = 0.5$; the assumption that σ is constant across households is identical to that in Auclert (2019). This figure shows

[1] Using the eleven paired capital gains yields in Bricker et al. (2025, Appendix A.3.2, Table 1, columns 6–7), I regress log gross private returns on log gross public returns with an intercept. The correlation between private and public returns is 0.3997. This calculation uses the published observations for 1989–2022, and is my calculation rather than an estimate reported by the authors.

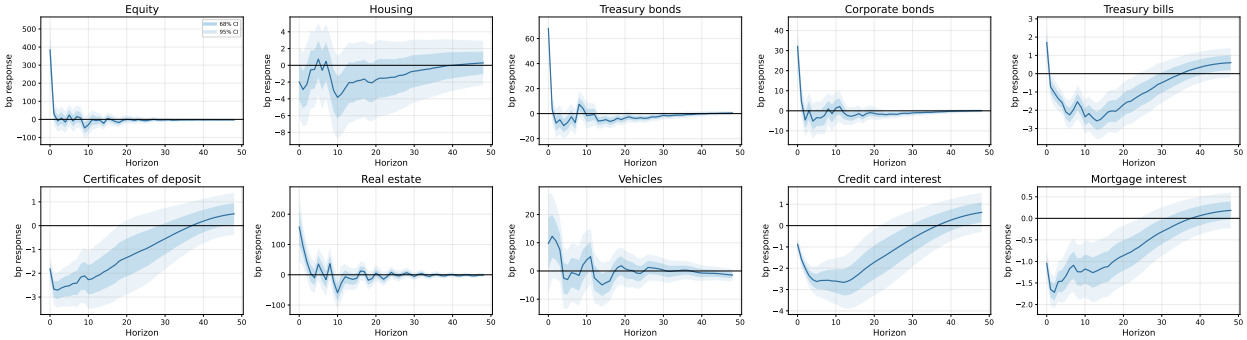


Figure 4: Impulse responses of asset returns to an expansionary monetary shock

the household's first year consumption response, as the estimated MPC is the annual change in consumption relative to an annual change in income. The first year consumption response is 0.286%, with a conditional 95% interval of $[0.122, 0.416]$%; revaluation contributes 0.280 percentage points. Standard errors are measured using 500 wild bootstrap draws.

The most striking channel in Figure 5 is the *revaluation* channel. That is, the change in household portfolio values during the first month of the monetary shock generates by far the greatest change in consumption. This is largely due to a sharp increase in equity prices following an expansionary monetary policy shock. The rate exposure channel is negative and significant, as a monetary expansion decreases future interest payments and hence discounted wealth. The substitution and income channels are positive but small, and the Fisher channel is not significant as the on-impact change in prices is small. Because the reference asset is the *wealth-weighted* aggregate portfolio, the first bars labeled "Aggregate portfolio reference" can be interpreted as the consumption response of the representative agent. In particular, the representative agent measurement would suggest a very large consumption response at about 0.4 percent following a 25 basis point decrease in the 2-year Treasury yield.

However, the portfolio composition channels *dampen* the aggregate effect. There is no statistically significant change in dividends or forward returns when considering heterogeneous portfolios, but the heterogeneous revaluation channel is *negative*. This implies a negative *sorting effect*: Low-MPC households hold assets that are particularly sensitive to the monetary shock, which dampens the total consumption response.

Because my estimation is at the household level, then aggregated, one advantage of this methodology is that I can measure the importance of arbitrary groups of households by regrouping households using different observables. Figure 6 shows the distribution of consumption across the wealth distribution, split into different channels for each wealth group. This figure shows a slight U-shape in the consumption response, reminiscent of the findings in Holm et al. (2021): Households roughly below median wealth have higher consumption responses than households between the 60-95th percentiles, while top-wealth households above the 95th percentile have very high consumption responses.

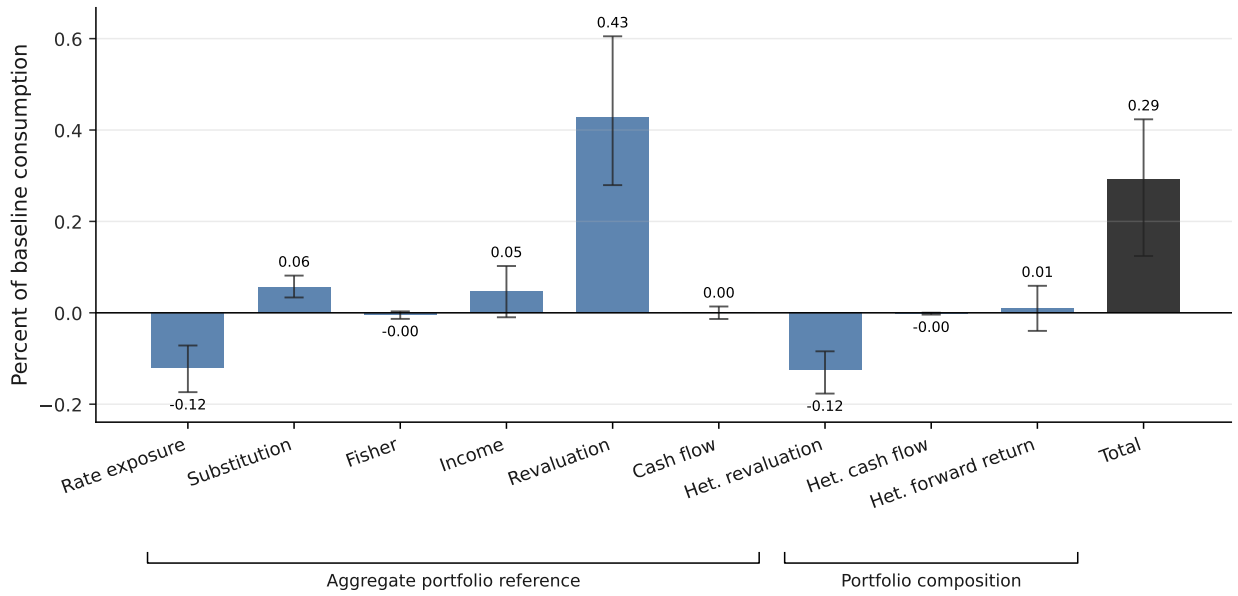


Figure 5: Decomposing the consumption response to expansionary monetary policy

Notes: Aggregate transmission channels following a 25 basis point decrease in the 2-year Treasury yield.

The primary transmission channels at the lower end of the wealth distribution are largely due to the income channel and the rate exposure channel, as lower rates also decrease debt payments going forward. At the top end of the wealth distribution, high-wealth households experience losses from lower future rates as their interest payments decrease, but this is greatly offset by an increase in the revaluation of their portfolios. In particular, because high-wealth households hold disproportionately more public and private equity than low-wealth households, the large positive equity price response to expansionary monetary policy has an outsize impact on their total budgets and consumption response. This is true *despite* the fact that high-wealth households tend to have lower MPCs, and that their MPC out of capital gains is even lower (e.g. 2.8% for the 95th-100th percentile in Di Maggio et al. (2020)).

A simple back-of-the-envelope calculation can directly show the importance of the concentration of equity holdings at the top of the wealth distribution. The public equity price response in the first month after the expansionary monetary shock in Figure 4 is 3.8%; including this response scaled by 0.4 for private equity, the average price response for broad equity is 2.7%. Households in the 95th-100th percentile hold 57% of their wealth in broad equity. Multiplying this by the baseline MPC out of capital gains for this group from Di Maggio et al. (2020) and taking the share over baseline aggregate consumption, the total consumption response for this group from equity revaluation only is:

$$\frac{\Delta C_{95\text{–}100}}{C} = \underbrace{0.028}_{\text{MPC}} \times \underbrace{0.027}_{\Delta Q/Q} \times \underbrace{0.57}_{equity/W} \times \underbrace{5.8}_{W/C_0} \approx 0.25\%.$$

This number is actually larger than the revaluation channel shown in Figure 6, because the formal normalization of the Di Maggio et al. (2020) MPC weights leads to a smaller MPC for many households that hold large equity shares. Thus, even though lower-wealth households have higher MPCs, especially out of income, these higher MPCs are not large enough to account for the strong equity response to monetary shocks. Thus, the revaluation channel dominates in the aggregate response of consumption to monetary policy, and this transmission is very concentrated in the top end of the wealth distribution.

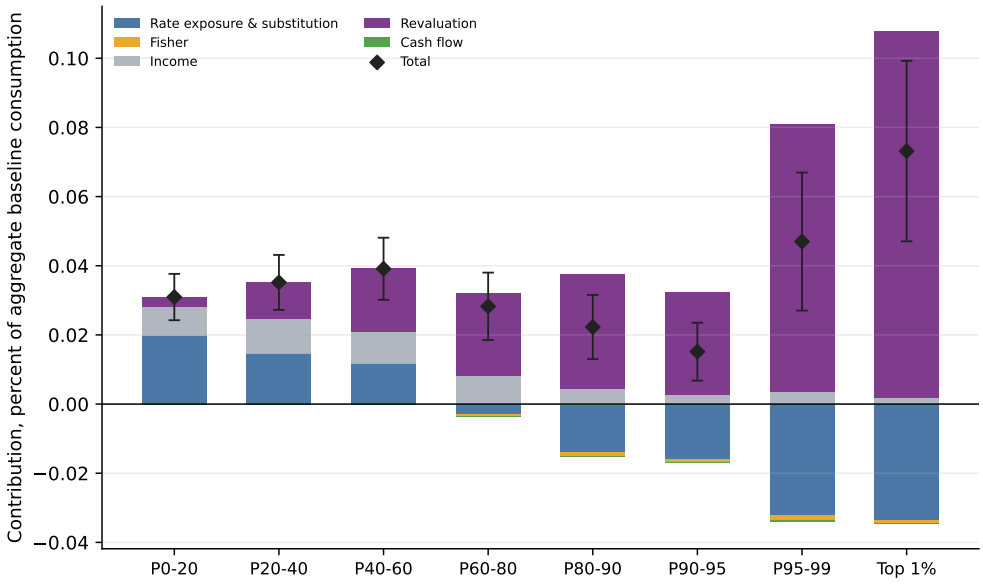


Figure 6: Transmission channels for an expansionary monetary shock across the wealth distribution

Notes: Distribution of transmission channels following a 25 basis point decrease in the 2-year Treasury yield.

3.6 Sensitivity to the revaluation MPC

To assess sensitivity to the calibration of the capital gain MPC in the main empirical results, I replace this revaluation MPC with a common coefficient, $MPC_i^Q = m$, across households and included priced assets. I use this common MPC to test two questions. First, what would the revaluation MPC have to be for the revaluation channel to be substantially smaller, such as the same magnitude as the income channel? Second, what would the revaluation MPC have to be for lower-wealth households to have a higher consumption response than higher-wealth households?

Figure 7 shows these two thresholds. In the top panel, I show the contribution of the revaluation channel to aggregate consumption in the solid blue line, where the baseline is the dashed blue horizontal line at 0.302. The baseline income channel is the dashed grey line. The revaluation MPC would have to be 0.0037, or 0.37 cents out of one dollar, in order for the revaluation channel to be the same size as the estimated income channel. This is two orders of magnitude smaller than common estimates in the literature (e.g. Di Maggio et al. (2020), Chodorow-Reich et al. (2021)),

which implies that the strength of the revaluation channel relative to the income channel is robust to reasonable estimates of the MPC out of capital gains.

In the bottom panel, I show the difference in the consumption response between the top wealth quintile (Q5) and the bottom wealth quintile (Q1). This difference is 0.127 in the baseline specification. In order for these groups to contribute the same consumption response, the MPC out of portfolio revaluation would have to be 0.01. This is still higher than most estimates in the literature, but is similar to the MPCs used in Slacalek et al. (2020). Appendix B.6 shows that under the same MPCs, my estimation methodology largely replicates the distribution of responses across poor hand-to-mouth (HtM), wealthy HtM, and non-HtM households in Slacalek et al. (2020).

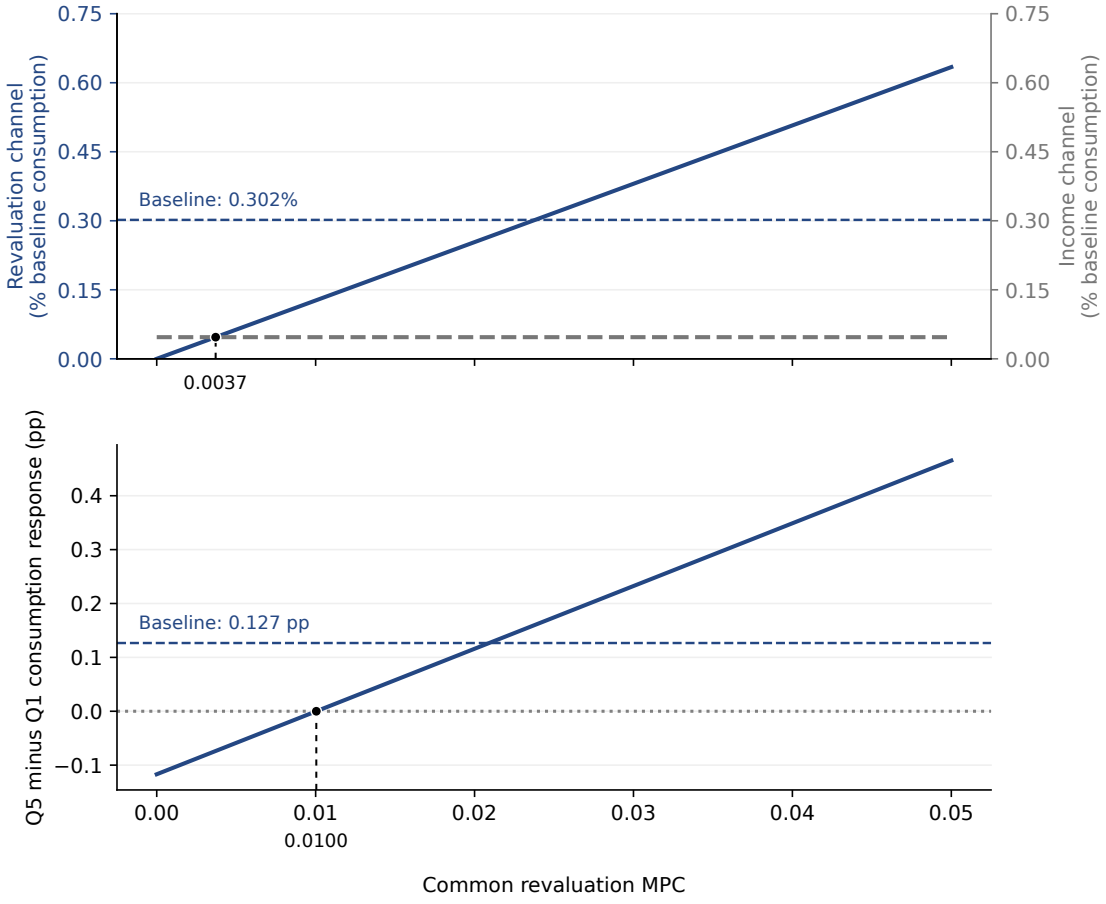


Figure 7: Sensitivity to a common revaluation MPC

Notes: The top panel reports the revaluation channel (left axis) and the income channel (right axis) as percentages of aggregate baseline consumption. The bottom panel reports the consumption response of the top minus bottom wealth quintile in percentage points of aggregate baseline consumption. Horizontal dashed blue lines show the calibrated baseline results; the grey dashed income line remains fixed at the baseline estimate. The MPC on the revaluation channel would have to be 0.003 for the revaluation channel to equal the income channel in magnitude, and 0.01 for the lowest wealth quintile to have a larger consumption response than the top wealth quintile.

3.7 Benchmarking to aggregate consumption responses

Next, I check the magnitude of the consumption response estimated in Section 3.5 in aggregate data from the Consumer Expenditure Survey. I use monthly data from the CEX from January 1990 through December 2023, and compute consumption following the procedure from Coibion et al. (2017). Consumption is winsorized at the 1st and 99th percentiles and deflated to 2019 dollars. I compute the impulse response of consumption to an expansionary monetary shock using the same method as in Section 3.4.

Figure 8 shows the impulse response of aggregate consumption to a monetary shock normalized to a 25 basis point decline in the two-year Treasury yield. Consumption at one year after the shock rises by approximately 0.4 percent, compared with the sufficient statistic estimate of 0.286 percent in Figure 5. Although a slight underestimate, the empirical decomposition estimate is well within the error bands of this aggregate IRF, which suggests that the decomposition methodology is relatively accurate in estimating the aggregate response. For further validation, Section 4 compares the empirical procedure with the consumption response in a full two-asset heterogeneous agent New Keynesian model and finds that the empirical decomposition is within single-digit percentage points of the full nonlinear solution.

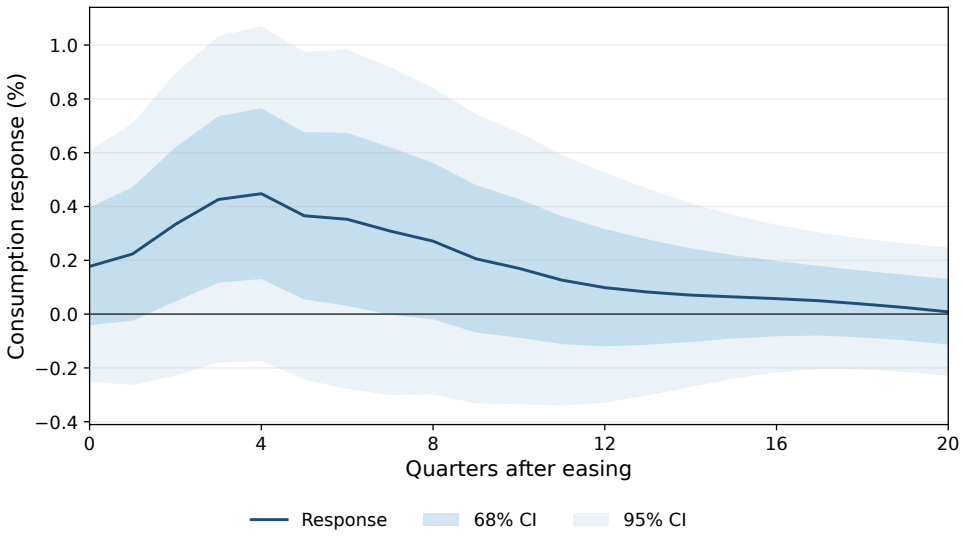


Figure 8: Aggregate consumption response following a monetary easing

Notes: Impulse response of aggregate consumption after a 25 basis point decrease in the 2-year Treasury yield.

Figure 9 shows the difference in the consumption response between wealth groups. The top row compares households in the highest and lowest deciles of financial asset holdings, while the bottom row compares households in the highest and lowest deciles of the share of financial assets held in equities. The left column shows these differences in dollars, while the right column shows these differences in percentages relative to each group's baseline consumption. The difference in consumption for the top financial asset decile is slightly larger than that of the bottom financial asset

decile on impact, although the difference is small and imprecisely estimated. However, households with the highest equity shares exhibit substantially larger responses than those with the lowest equity shares. This pattern is consistent with the main mechanism from the empirical decomposition, that portfolio composition and especially equity holdings matter for monetary transmission.

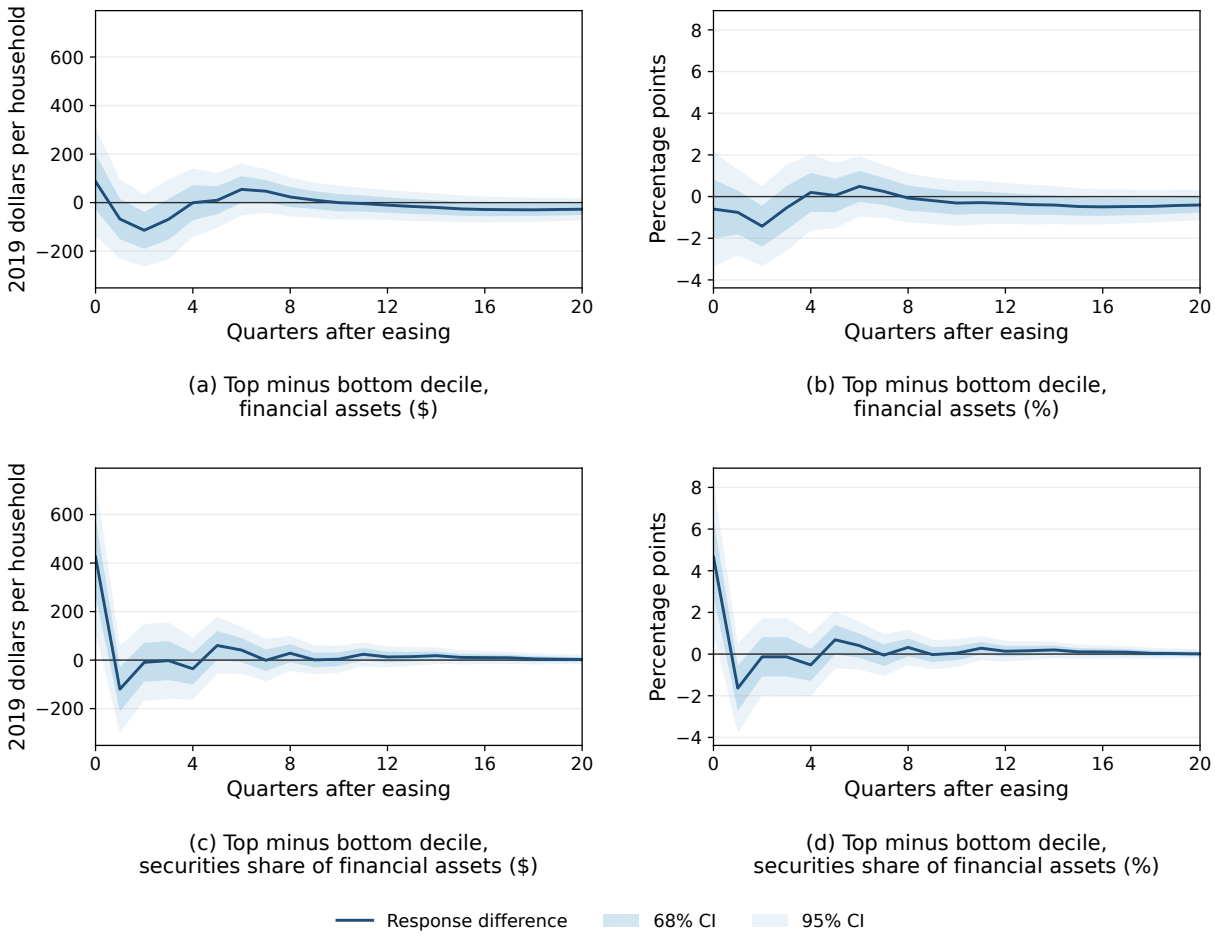


Figure 9: Differences in consumption responses across groups following a monetary easing

4 Quantitative validation and the distribution of responses

In this section, I use a quantitative New Keynesian model to validate the first-order consumption decomposition and study the distribution of consumption responses in the model. I compare the path and distribution of responses to a monetary innovation as well as a debt-financed government spending shock in both a canonical two-asset HANK model as well as an extension that matches moments of the wealth distribution and equity price responses in the data.

4.1 Model structure overview

In this section, I study two versions of a two-asset HANK model: The canonical two-asset HANK model from Auclert et al. (2021) as well as an extended model that adds household wealth and equity exposure types calibrated to the data. I start with the two-asset model from Auclert et al. (2021) to compare the first-order consumption decomposition with a full nonlinear response. Then, I compare the distribution of responses between these two models, as well as versions of these models that are augmented with a wedge on equity returns to better match the impulse response of equity prices to monetary shocks.

In the canonical model from Auclert et al. (2021), households hold a liquid asset b_{it}, an illiquid asset a_{it}, and face convex portfolio adjustment costs $\Phi_t(a_{it}, a_{it-1})$. A financial intermediary takes household deposits and returns a liquid and illiquid asset subject to a liquidity transformation cost, which leads to a higher return on the illiquid asset. The household and financial blocks, together with the model's accounting identities, are described in Appendix C. The only modification I make to this model is to calibrate the aggregate firm equity to net worth ratio to match the empirical share of broad equity at 44.7%, compared with 80 percent in the canonical calibration. The remaining assets are government bonds and firm debt. As in the canonical model, liquid assets are composed of government bonds only. Illiquid assets are provided by the financial intermediary and include government bonds, firm equity, and firm debt.

Previous work has documented that HANK models struggle to match moments of the wealth distribution (e.g. Kaplan et al. (2018), Benhabib and Bisin (2018)) and equity price responses to shocks (e.g. Kaplan and Violante (2018), Kekre and Lenel (2022)). However, the empirical results from Section 3 emphasize the unequal incidence of monetary transmission across the wealth distribution, partly due to a large equity price response. To study a model that matches the equity price response and its impact across the wealth distribution, I incorporate two extensions to the canonical two-asset HANK model. First, I calibrate the equity price response in the benchmark model to the empirical impulse response of prices to a monetary shock, motivated by the inelastic intermediary portfolio demand argument from Gabaix and Koijen (2021). I complement this benchmark model with an extended version that assigns households to *permanent wealth types* that match the wealth and equity holding distribution in the data. In the following, I compare four different models: The canonical model and the wealth-types model, both including and without the equity valuation wedge.

4.2 Calibration: Wealth types and equity valuation

Wealth types I extend the benchmark two-asset model of Auclert et al. (2021) with five permanent types of equal population mass, differing in patience β_j and equity exposure θ_j. I fit each quintile's share of total net worth and of aggregate broad equity to match the averages in the pooled 1989–2025 SCF, where wealth quintiles are defined within each wave. I define broad equity as the

sum of public equity holdings and private business wealth. Figure 10 and Appendix C.2 report the targets and achieved moments using the same definitions and sample.

Equity valuation after a shock To better match the impulse response of equity prices following a monetary shock, I modify the financial intermediary's problem to include a portfolio constraint on equity holdings, which can be calibrated to generate greater equity volatility as in Gabaix and Koijen (2021). Under perfect foresight, the financial intermediary chooses its own equity share θ_t^I to maximize returns subject to portfolio constraints:

$$\max_{\underline{\theta}_t^I \leq \theta_t^I \leq \bar{\theta}_t^I} \left\{ \theta_t^I \frac{d_{t+1} + p_{t+1}}{p_t} + (1 - \theta_t^I)(1 + r_{t+1}) \right\} \tag{21}$$

where p_t is the equity price, d_t is the dividend, and r_{t+1} is the safe real interest rate. The intermediary takes these values as well as its portfolio constraints $(\underline{\theta}_t^I, \bar{\theta}_t^I)$ as given.

Let μ_t^+ and μ_t^- denote the Lagrange multipliers on the upper and lower portfolio constraints, respectively. The intermediary's first-order condition implies:

$$\begin{aligned} &\frac{d_{t+1} + p_{t+1}}{p_t} = 1 + r_{t+1} + \underbrace{\mu_t^+ - \mu_t^-}_{\zeta_t}, \\ &\mu_t^+(\bar{\theta}_t^I - \theta_t^I) = 0, \qquad \mu_t^-(\theta_t^I - \underline{\theta}_t^I) = 0 \end{aligned} \tag{22}$$

where ζ_t is the *equity valuation wedge* resulting from the intermediary's portfolio constraints. The upper bound limits the intermediary's equity investment when equity offers a higher return than bonds; the lower bound guarantees the intermediary's equity investment when equity offers a lower return than bonds. These constraints therefore allow ζ_t to take both positive and negative values and drive an additional wedge between the equity return and the risk-free rate.

The broad equity asset in the model represents the sum of public equity and private business wealth. To accurately capture the price response, I target the average price response of broad equity in the data weighted by the share of wealth held in each asset. Specifically, I calibrate ζ_t in each economy to match the impact and half-life of the wealth-weighted nominal public/private equity price response as measured in Section 3. The top right panel of Figure 10 shows the results of this calibration. I make the following functional form assumption on the ζ_t wedge:

$$\zeta_t = \rho_\zeta \zeta_{t-1} + \chi_\zeta \Delta r_t^*, \qquad \zeta_{-1} = 0 \tag{23}$$

Here, χ_ζ is the sensitivity of the valuation wedge to the change in the Taylor rule intercept Δr_t^* and ρ_ζ is the persistence of the response. Households receive returns on their portfolios depending

on their type's equity exposure θ_j

$$\begin{aligned} 1 + r^a_{jt} &= \theta_{jt} \frac{d_t + p_t}{p_{t-1}} + (1 - \theta_{jt})(1 + r_t) \\ &= \theta_{jt}(1 + r_t + \zeta_{t-1}) + (1 - \theta_{jt})(1 + r_t) \\ &= 1 + r_t + \theta_{jt}\zeta_{t-1}, \qquad t \geq 1, \qquad r^b_t = r_t - \omega. \end{aligned} \tag{24}$$

where ω is the liquidity transformation cost paid by the intermediary. Here θ_{jt} is type j's equity share chosen at $t-1$, which may vary over time to ensure that the aggregate equity share remains constant. Appendix C.2 further explains how equity markets clear and how the date-0 price change affects household wealth.

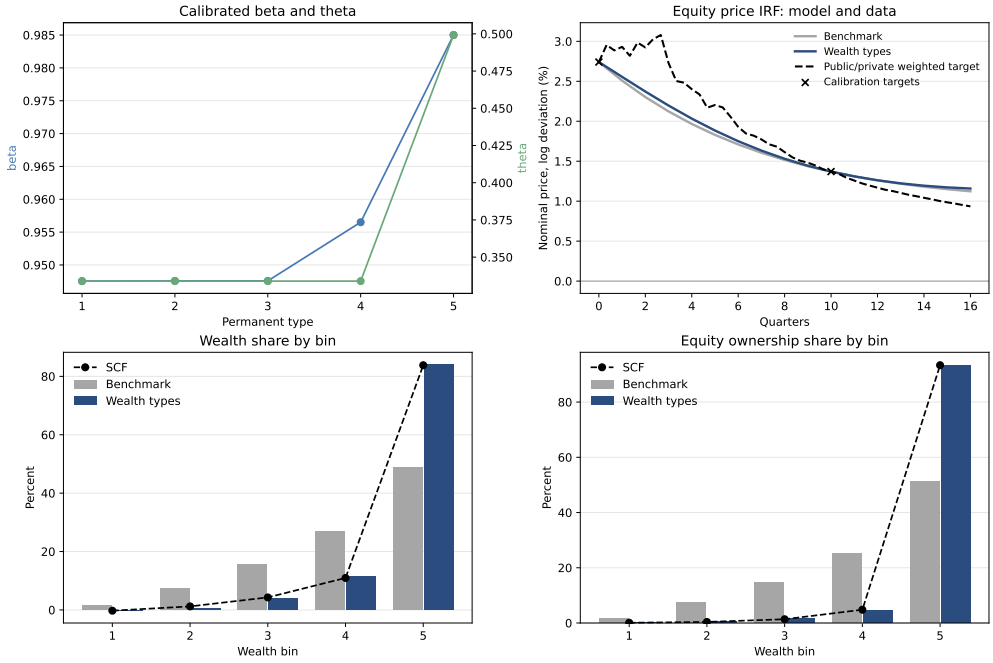


Figure 10: Calibration of wealth moments and equity prices

4.3 Validation of sufficient statistics against the nonlinear model

The first exercise I conduct in the quantitative model is to validate the sufficient statistics decomposition in an environment where the full nonlinear response is known. Because the decomposition in Section 2 is a first-order approximation, it does not necessarily provide the exact same consumption response as in the full nonlinear model. Additionally, the empirical procedure maintains fixed initial portfolio exposures and steady-state MPCs. These MPC estimates better approximate the empirical procedure, and incorporate local saving and adjustment responses but are not direct Jacobians estimated along the transition path. The comparison can therefore contain errors from the sufficient statistics mapping, household nonlinearities, and any unmatched anticipated news.

Appendix C.4 separates these terms.

Within each model, I compute the consumption response to a 25 quarterly basis point monetary innovation in the first year following the shock using both the empirical methodology as well as the nonlinear solution. To replicate the empirical methodology, I compute MPCs out of cash, capital gains, and income by measuring derivatives of expected consumption at the steady states. In the nonlinear computation, I solve the model using the usual procedure: Solve household policies backward and the household distribution forward until markets clear. The empirical procedure retains news from the first year after the shock. I compare it with the nonlinear household solution using the same news and in the full equilibrium, which retains the complete anticipated path.

Figure 11 shows the comparison between the empirical procedure and full nonlinear equilibrium. The empirical line uses news through the first year, whereas the full solution retains the full path of future information. For the 25 quarterly basis point easing with the additional equity wedge (bottom row), the gap between the empirical methodology and the nonlinear solution is 0.0688 percentage points of annual baseline consumption in the benchmark model and 0.0606 in the wealth type model. When comparing the nonlinear solution with an identical news truncation at one year, the empirical procedure exceeds the household Jacobian by only 0.0133 and 0.0123 percentage points, respectively. The gaps between the models without the equity wedge are smaller, as the large equity wedge adds nonlinearities. Appendix C.4 reports the numerical differences and separates between errors from nonlinearities and from the time truncation. Although the difference between the empirical methodology and the full nonlinear solution increases with the shock size as nonlinearities become more prominent, the empirical methodology is a fairly good approximation of the full nonlinear response at reasonable shock sizes in all models.

4.4 Model-implied responses to a monetary innovation

Next, I compare the distribution of wealth and consumption changes across households, as well as the aggregate impulse response of aggregate variables, in response to an expansionary monetary policy shock. To do so, I compare the benchmark and wealth types models, both with and without the equity valuation wedge. In all models, I impose the same monetary innovation under the same Taylor rule:

$$\begin{aligned} \Delta i_t &= \Delta r_t^* + \phi_\pi \Delta \pi_t, \qquad \phi_\pi = 1.5, \\ \Delta r_t^* &= -0.0025(0.8)^t. \end{aligned} \tag{25}$$

The shock is a 25 basis point reduction in the quarterly nominal rate intercept on impact, which decays at rate 0.8. Both models retain the 44.7 percent aggregate broad-equity share described above. In each model, the wedge follows Equation (23) with its previously calibrated parameters.

Figure 12 shows the impulse responses of consumption, output, and inflation, as well as the real and nominal rate and the nominal equity price, following this monetary shock. The equity

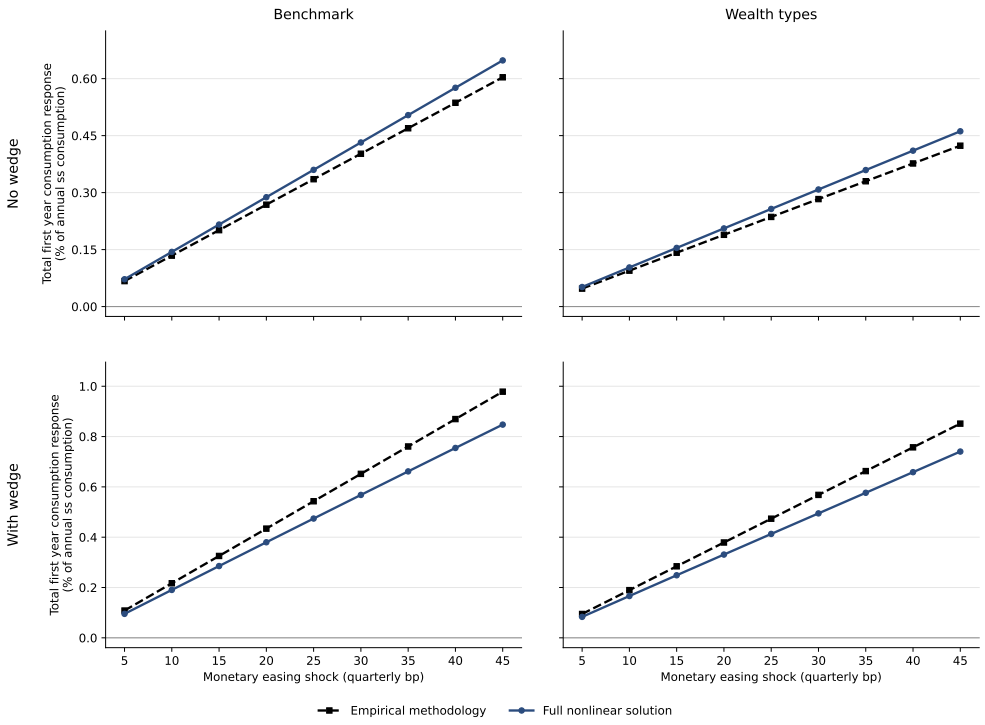


Figure 11: Consumption change: Empirical methodology versus nonlinear solution

Notes: The empirical methodology, in solid blue, is generally within single digit percentage points of the full nonlinear solution. This suggests that the empirical methodology is a decent approximation of the full nonlinear response in the two-asset HANK models used.

valuation wedge increases first-year consumption from 0.360 to 0.477 percent in the benchmark model and from 0.258 to 0.417 percent in the wealth types model. Thus, matching the larger equity price response in the data amplifies consumption transmission in both models, consistent with our empirical finding that equity revaluation is a key transmission channel of monetary policy. However, the wealth types model generates a smaller consumption response than the benchmark both with and without the wedge. This is also consistent with our empirical finding: The *sorting* of low-MPC, high-wealth households into larger holdings of equity *dampens* the consumption effect of monetary policy.

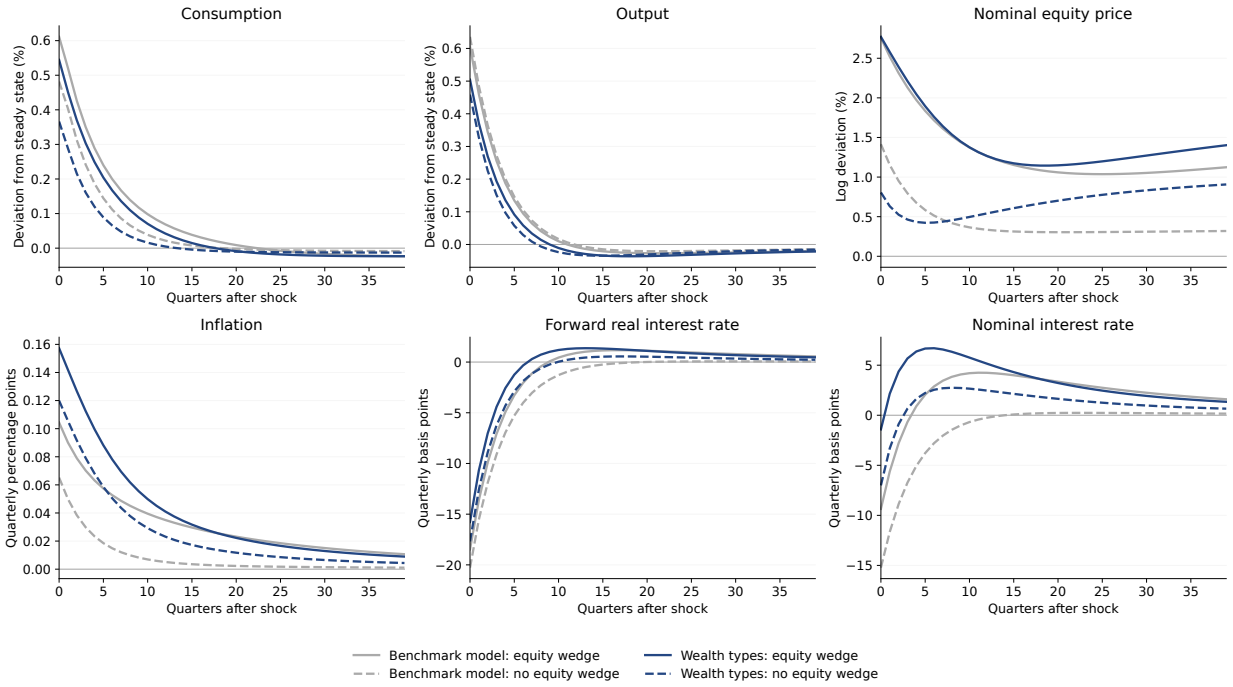


Figure 12: Responses to a monetary easing under the Taylor rule

Notes: Impulse responses across different models to the monetary innovation in Equation (25).

Figure 13 compares wealth gains and consumption responses across household wealth quintiles, defined in steady state before the shock. The left panel shows the share of total revaluation gains allocated to each wealth quintile, while the right panel shows the share of the aggregate consumption change allocated to each wealth quintile. The wealth-types model better matches the distribution of both wealth gains and consumption changes following the shock. In the left panel, both wealth-types models almost exactly match the empirical share of revaluation gains, with the top wealth quintile receiving the vast majority of these gains. This is largely due to top wealth quintiles holding more initial wealth and thus being more exposed to the shock. Additionally, the wealth-types model with the equity wedge better matches the increase in consumption for all quintiles of the wealth distribution compared to the other models. This is both because of the better match to the distribution of revaluation gains, as well as the level as the equity price wedge increases the total amount of revaluation gains especially for those in the highest wealth quintile.

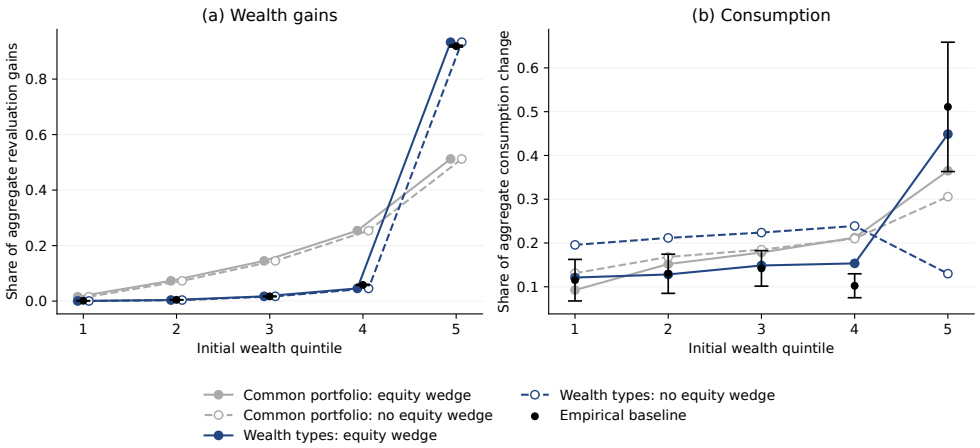


Figure 13: Shares of revaluation gains and consumption changes by wealth quintile

Notes: Adding all points from a model sums to one, which is the normalization of the aggregate response. The wealth-type plus equity wedge model better matches both the distribution of wealth gains and consumption changes after an expansionary monetary shock compared with the benchmark model.

One key takeaway from these comparisons is that better matching the empirical distribution of wealth and equity prices substantially changes the analysis of distributional consequences of monetary policy. Recent work such as McKay and Wolf (2023) has suggested that the distribution of consumption after a monetary shock is relatively flat across the wealth distribution. However, both the empirical results in Section 3 as well as the quantitative results from the wealth-types model with the equity wedge find that monetary policy is highly *regressive*, with consumption gains from an expansionary shock concentrated among the wealthiest households due to their large equity holdings.

4.5 Fiscal policy and consumption across the wealth distribution

Next, I study how household portfolios affect the distribution of consumption responses to a government spending shock. Because the equity price wedge is not calibrated to respond to a government spending shock, differences between the benchmark and wealth-types model will be due solely to the additional wealth parameters. I retain the previously specified calibrations as well as the Taylor coefficient of $\phi_\pi = 1.5$, but do not incorporate an additional monetary innovation. The experiment is a one percent increase in government spending as a percent of steady-state quarterly output, which decays at rate 0.8:

$$G_t - \bar{G} = 0.01\bar{Y}(0.8)^t, \qquad t \geq 0. \tag{26}$$

This government spending shock is *debt-financed* to capture intertemporal effects and deviations from Ricardian equivalence in both economies. To fund greater spending, the government issues

debt held in household accounts, and delays changes in the labor tax rate for four quarters:

$$B_t^g = (1 + r_t)B_{t-1}^g + G_t - \tau_t w_t N_t,$$
$$\tau_t - \bar{\tau} = \begin{cases} 0, & t = 0, \ldots, 3, \\ 0.05(B_{t-1}^g - \bar{B}^g)/(\bar{w}\bar{N}), & t \geq 4. \end{cases} \tag{27}$$

Households learn the full path of government spending and tax rates when the shock occurs. Note that holding the tax rate fixed initially does not hold tax revenues fixed: Earnings can change the tax base, while inflation changes the real cost of servicing inherited debt. Debt therefore does not necessarily rise one for one with spending.

Figure 14 plots consumption responses by initial wealth quintile. I follow the same households through the transition and compare their consumption with and without the government spending shock. The aggregate change in consumption in the first year is identical, at -0.119 percent in both models. A fall in consumption following the increase in government spending is due to the classic crowding-out effect of monetary policy, similar to findings in Hagedorn et al. (2019), Ferriere and Navarro (2025), and others. However, the distribution of this effect is heterogeneous across households. The consumption response by the lowest quintile households in the first year is -0.079 in the canonical model and 0.142 percent in the wealth-types model. At the top, the responses are -0.126 and -0.200 percent, respectively.

Thus, this shock is much more *progressive* in the wealth-type model compared with the canonical model. This is largely because the lowest wealth quintile households in the wealth-type model are net debtors in government bonds and hold almost no equity, so the inflationary effect of government spending reduces their real debt burden. In contrast, the lowest quintile in the canonical model holds positive safe assets and more equity, so the change in equity payoffs reduces their wealth and therefore consumption. However, these immediate differences are somewhat temporary; consumption for the lowest wealth quintile in the wealth-type model eventually falls when the tax rate adjusts.

Overall, this fiscal experiment shows that similar aggregate consumption responses to shocks such as a government spending shock can conceal different distributions and timing of gains and losses across the distribution of households. This illustrates the interaction between household wealth and portfolio exposures and policy, and suggests that incorporating more realistic household wealth characteristics into quantitative models may change the policy evaluations.

5 Conclusion

This paper studies how monetary policy affects households through changes in asset prices and dividends, which flow through the economy depending on household-specific portfolios. I derive a sufficient statistics decomposition of the consumption response to a monetary policy shock that de-

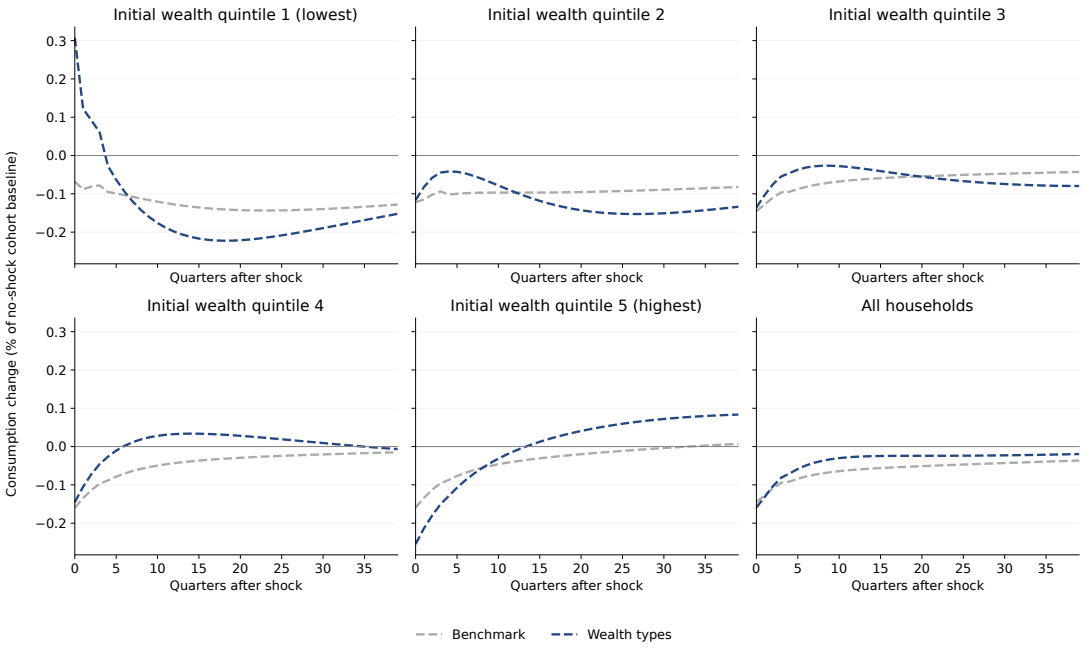


Figure 14: Consumption responses to government spending by initial wealth quintile

Notes: Both models use the spending shock and delayed-tax rule in Equations (26)–(27), with Taylor coefficient $\phi_\pi = 1.5$ and no equity wedge. Each initial quintile contains 20 percent of its model's population, with tied wealth holdings split proportionally. Responses are percentages of each fixed cohort's consumption at the same date without the shock, including changes in both consumption policies and cohort distributions. The figure shows the first 40 quarters of a 900-quarter first-order solution.

pends on the joint distribution of household portfolio exposures, marginal propensities to consume, and changes in asset returns and payoffs following a monetary shock. I measure these objects using public U.S. data and find that the portfolio revaluation channel is a particularly strong channel for monetary policy transmission. This channel is largely driven by a large equity price response to monetary shocks, which generates substantial wealth gains for high-wealth households.

A two-asset HANK model with permanent wealth types and an equity valuation wedge better matches the data than a canonical HANK model. Specifically, the extended model allocates a large share of the wealth gains and consumption response to households at the top of the wealth distribution. These results highlight the importance of jointly matching household portfolio exposures as well as asset price responses when studying the distributional effects of monetary policy transmission. As in the data, the extended model exhibits an amplifying effect of equity prices and a dampening effect of portfolio sorting. A government spending experiment further shows that similar aggregate consumption responses may conceal different distributional responses depending on the model calibration. Future work could shed further light on these portfolio channels by using more detailed microdata to study how households adjust their portfolios in response to shocks, and by examining how the dynamics of asset prices following monetary shocks arise in general equilibrium.

References

Adrian, Tobias, Richard K Crump, and Emanuel Moench, "Pricing the term structure with linear regressions," *J. Financ. Econ.*, October 2013, *110* (1), 110–138. 11, 51

Amberg, Niklas, Thomas Jansson, Mathias Klein, and Anna Rogantini Picco, "Five facts about the distributional income effects of monetary policy shocks," *Am. Econ. Rev. Insights*, September 2022, *4* (3), 289–304. 5

Andersen, Asger Lau, Niels Johannesen, Mia Jørgensen, and José-Luis Peydró, "Monetary policy and inequality," *J. Finance*, July 2023, *78* (5), 2945–2989. 5

Auclert, Adrien, "Monetary policy and the redistribution channel," *Am. Econ. Rev.*, June 2019, *109* (6), 2333–2367. 4, 5, 6, 7, 10, 13, 14, 16, 23, 51

_ , Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, September 2021, *89* (5), 2375–2408. 3, 30, 54

Bach, Laurent, Laurent E Calvet, and Paolo Sodini, "Rich pickings? Risk, return, and skill in household wealth," *Am. Econ. Rev.*, September 2020, *110* (9), 2703–2747. 5

Bauer, Michael D and Eric T Swanson, "A reassessment of monetary policy surprises and high-frequency identification," *NBER Macroecon. Annu.*, May 2023, *37*, 87–155. 22, 23

Becko, John Sturm, Joel P. Flynn, and Christina Patterson, "Fiscal Policy in a Networked Economy," June 2026. Working paper, June 20. 21

Benhabib, Jess and Alberto Bisin, "Skewed wealth distributions: Theory and empirics," *J. Econ. Lit.*, December 2018, *56* (4), 1261–1291. 30

Bernanke, Ben S and Kenneth N Kuttner, "What explains the stock market's reaction to Federal reserve policy?," *J. Finance*, June 2005, *60* (3), 1221–1257. 2, 23

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *Am. Econ. Rev.*, 2008, *98* (5), 1887–1921. 16, 20, 48, 49

Bricker, Jesse, Alice Henriques, and Kevin Moore, "Updates to the Sampling of Wealthy Families in the Survey of Consumer Finances," Finance and Economics Discussion Series 2017-114, Board of Governors of the Federal Reserve System, Washington, DC 2017. 16

_ , Kevin B. Moore, and Alice H. Volz, "Rates of Return on Private and Public Businesses," Finance and Economics Discussion Series 2025-107, Board of Governors of the Federal Reserve System, Washington, DC 2025. 23

Canto, Felipe Del, John Grigsby, Eric Qian, and Conor Walsh, "Are Inflationary Shocks Regressive? A Feasible Set Approach," *Q. J. Econ.*, 2024. 4

Chodorow-Reich, Gabriel, Plamen T Nenov, and Alp Simsek, "Stock market wealth and the real economy: A local labor market approach," *Am. Econ. Rev.*, May 2021, *111* (5), 1613–1657. 18, 26, 54

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia, "Innocent Bystanders? Monetary policy and inequality," *J. Monet. Econ.*, June 2017, *88*, 70–89. 28

Doepke, Matthias and Martin Schneider, "Inflation and the Redistribution of Nominal Wealth," *J. Polit. Econ.*, December 2006, *114* (6), 1069–1097. 4

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, "Heterogeneity and persistence in returns to wealth," *Econometrica*, 2020, *88* (1), 115–170. 5, 16, 17

Fernández-Villaverde, Jesús and Oren Levintal, "The distributional effects of asset returns," Technical Report w32182, National Bureau of Economic Research, Cambridge, MA March 2024. 5

Ferriere, Axelle and Gaston Navarro, "The heterogeneous effects of government spending: It's all about taxes," *Rev. Econ. Stud.*, March 2025, *92* (2), 1061–1125. 37

Gabaix, Xavier and Ralph S. J. Koijen, "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis," Working Paper 28967, National Bureau of Economic Research June 2021. 3, 30, 31

Gilchrist, Simon and Egon Zakrajšek, "Credit spreads and business cycle fluctuations," *Am. Econ. Rev.*, June 2012, *102* (4), 1692–1720. 22

Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman, "The Fiscal Multiplier," Technical Report w25571, National Bureau of Economic Research, Cambridge, MA February 2019. 37

Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek, "The transmission of monetary policy under the microscope," *J. Polit. Econ.*, October 2021, *129* (10), 2861–2904. 5, 24

Kang-Landsberg, Alena and Matthew Plosser, "How Do Deposit Rates Respond to Monetary Policy?," *Liberty Street Economics*, 2022, (November 21). Liberty Street Economics Blog. 2

Kaplan, Greg and Giovanni L Violante, "Microeconomic Heterogeneity and Macroeconomic Shocks," *J. Econ. Perspect.*, 2018, *32* (3), 167–194. 30

_ , Benjamin Moll, and Giovanni L Violante, "Monetary policy according to HANK," *Am. Econ. Rev.*, March 2018, *108* (3), 697–743. 4, 30

Kekre, Rohan and Moritz Lenel, "Monetary policy, redistribution, and risk premia," *Econometrica*, September 2022, *90* (5), 2249–2282. 4, 30

Lee, Donggyu, "Unconventional Monetary Policies and Inequality," Staff Reports 1108, Federal Reserve Bank of New York July 2024. Revised October 2026. 5

Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S Mitchell, "Optimal financial knowledge and wealth inequality," *J. Polit. Econ.*, April 2017, *125* (2), 431–477. 6

Maggio, Marco Di, Amir Kermani, and Kaveh Majlesi, "Stock market returns and consumption," *J. Finance*, December 2020, *75* (6), 3175–3219. 18, 21, 25, 26, 54, 57

McKay, Alisdair and Christian K Wolf, "Monetary policy and inequality," *J. Econ. Perspect.*, February 2023, *37* (1), 121–144. 36

Melcangi, Davide and Vincent Sterk, "Stock market participation, inequality, and monetary policy," *Rev. Econ. Stud.*, June 2024, *92* (4), 2656–2690. 4

Menzio, Guido and Saverio Spinella, "A Quantitative Theory of Heterogeneous Returns to Wealth," Technical Report w33868, National Bureau of Economic Research June 2025. 6

Patterson, Christina, "The matching multiplier and the amplification of recessions," *Am. Econ. Rev.*, April 2023, *113* (4), 982–1012. 3, 16, 17, 18, 19, 21, 48

Slacalek, Jiri, Oreste Tristani, and Giovanni L Violante, "Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area," *J. Econ. Dyn. Control*, June 2020, *115* (103879), 103879. 4, 27, 54

Vissing-Jorgensen, Annette, "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," *NBER Work. Pap. Ser.*, April 2022, (8884). 6

A Household theory and proofs

A.1 Optimality conditions for the household's problem

Under the fixed portfolio assumption $\mathbf{a}_{it} = \boldsymbol{\omega}_i a_{it}$ in Equation (1), the household's first-order conditions are:

$$\lambda_{it} = \frac{u'(c_{it})}{P_t}, \qquad v'(n_{it}) = \lambda_{it} W_{it},$$
$$\lambda_{it} Q_{it} = \beta_i \lambda_{it+1}(Q_{it+1} + D_{it+1}), \qquad t < T,$$

where λ_{it} is the multiplier on the nominal budget constraint. The household's Euler equation is thus $u'(c_{it}) = \beta_i R_{it+1} u'(c_{it+1})$, where R_{it+1} is idiosyncratic and depends on the household's portfolio shares. The terminal condition is $a_{i,T+1} = 0$. Appendix A.5 gives the corresponding conditions with portfolio constraints that generate the portfolio shares.

A.2 Present-value budget and initial revaluation

Let $q_{it} = 1/\mathcal{R}_{it}$, with $q_{i0} = 1$. Multiplying the household's budget constraint by q_{it}, using $q_{it} R_{it} = q_{i,t-1}$, and imposing the terminal condition gives:

$$\sum_{t=0}^{T} q_{it} c_{it} = A_{i0} + \sum_{t=0}^{T} q_{it} Y_{it} = \mathcal{W}_{i0}. \tag{A.1}$$

With unrestricted saving and borrowing, this present value budget constraint is locally equivalent to the sequence of time t budget constraints. This object is helpful in working with household wealth terms. Conditional on the path of labor, separability between consumption and labor reduces the household problem to maximizing $\sum_{t=0}^{T} \beta_i^t u(c_{it})$ subject to (A.1).

Initial revaluation (Lemma 1). Differentiating $A_{i0} = P_0^{-1} \sum_k (Q_0^k + D_0^k) a_{i0}^k$ at the household's predetermined holdings gives

$$dA_{i0} = P_0^{-1} \sum_k a_{i0}^k (dQ_0^k + dD_0^k) - A_{i0}\, d\log P_0.$$

Thus, the change in a household's portfolio value is independent of subsequent portfolio choices.

A.3 Proof of Lemma 3

The household's first-order conditions are $\beta_i^t u'(c_{it}) = \lambda_i q_{it}$ with $\lambda_i > 0$. Define the household's intertemporal elasticity of substitution (IES) σ_{it} as well as the weighted present value of the house-

hold's lifetime consumption B_i as:

$$\sigma_{it} = -\frac{u'(c_{it})}{c_{it}u''(c_{it})} > 0, \qquad B_i = \sum_{t=0}^{T} q_{it}\sigma_{it}c_{it} > 0.$$

Differentiating gives

$$dc_{it} = -\sigma_{it}c_{it}(d\log\lambda_i + d\log q_{it}). \tag{A.2}$$

At fixed intertemporal prices, the budget implies $d\mathcal{W}_{i0} = -B_i\, d\log\lambda_i$, so

$$MPC_i = \frac{\partial c_{i0}}{\partial \mathcal{W}_{i0}} = \frac{\sigma_{i0}c_{i0}}{B_i}.$$

For a compensated change in the future path of interest rates, fixed consumption utility implies $\sum_t q_{it}dc_{it} = 0$, so

$$d\log\lambda_i = -\frac{\sum_t q_{it}\sigma_{it}c_{it}\, d\log q_{it}}{B_i}.$$

Using $q_{i0} = 1$ and $d\log q_{it} = -d\log\mathcal{R}_{it}$ yields

$$\varepsilon^h_{i0t} = -\frac{\partial\log c^h_{i0}}{\partial\log\mathcal{R}_{it}} = \frac{\sigma_{i0}q_{it}\sigma_{it}c_{it}}{B_i}, \qquad \sum_{t=1}^{T}\varepsilon^h_{i0t} = \sigma_{i0}(1 - MPC_i).$$

For $T \geq 1$, $0 < MPC_i < 1$. Under time-consistent IES coefficients in Assumption 6, factoring out the IES gives the weights in (20); the IES need not be common across households. For one-period returns, the corresponding coefficients are tail sums:

$$\sum_{t=1}^{T}\varepsilon^h_{i0t}\, d\log\mathcal{R}_{it} = \sum_{s=1}^{T}\left(\sum_{t=s}^{T}\varepsilon^h_{i0t}\right) d\log R_{is}.$$

A.4 Proof of Proposition 1

Differentiating (A.1) and subtracting the change in the future path of planned consumption gives $d\Omega_{i0}$, or the change in lifetime wealth after accounting for changes in the cost of the household's consumption plan:

$$\begin{aligned} d\Omega_{i0} &\equiv d\mathcal{W}_{i0} - \sum_{t=1}^{T} c_{it}\, dq_{it} \\ &= dA_{i0} + \sum_{t=0}^{T} q_{it}\, dY_{it} + \sum_{t=1}^{T} q_{it}(c_{it} - Y_{it})\, d\log\mathcal{R}_{it}. \end{aligned} \tag{A.3}$$

Since $\sum_t q_{it} dc_{it} = d\Omega_{i0}$, Equation (A.2) implies

$$-B_i \, d\log \lambda_i = d\Omega_{i0} + \sum_{t=1}^{T} q_{it} \sigma_{it} c_{it} \, d\log q_{it}.$$

Multiplying by $\sigma_{i0} c_{i0} / B_i$ gives

$$dc_{i0} = MPC_i d\Omega_{i0} - c_{i0} \sum_{t=1}^{T} \varepsilon^h_{i0t} \, d\log \mathcal{R}_{it}.$$

Substituting in for the previously derived $d\Omega_{i0}$ gives the desired (7).

MPC derivation Holding the path of labor and intertemporal prices fixed, the change in the household's total wealth due to the shock in Definition 1 is $d\mathcal{W}_{i0} = dMV_{i0} + dCF_{i0} + \sum_t q_{it} dY_{it}$. Thus the change in consumption resulting from each channel is:

$$\frac{\partial c_{i0}}{\partial MV_{i0}} = \frac{\partial c_{i0}}{\partial CF_{i0}} = \frac{\partial c_{i0}}{\partial Y_{i0}} = MPC_i, \qquad \frac{\partial c_{i0}}{\partial Y_{it}} = MPC_i q_{it}.$$

This overall consumption decomposition allows labor supply to respond, where changes in earnings will be included in dY_{it}, while MPC_i remains conditional on the labor choice.

A.5 Additional notation and proof of Lemma 2

At each date $t < T$, suppose there is locally unrestricted borrowing and saving given some portfolio shares $\omega \in F_{it}$, where F_{it} is independent of expenditure, wealth, returns, and the shock. Define the household's maximal real return as

$$R^*_{i,t+1}(\epsilon) = \max_{\omega \in F_{it}} \frac{[Q_{t+1}(\epsilon) + D_{t+1}(\epsilon)]'\omega}{Q_t(\epsilon)'\omega} \frac{P_t(\epsilon)}{P_{t+1}(\epsilon)}.$$

Assume that this maximum exists and is differentiable at zero. Additionally, assume that some common portfolio composition $\bar{\omega}_i$ attains this maximum at the baseline at every date. Then, both the constrained and unconstrained problems share the same portfolio compositions, preferences, paths of future inputs, and terminal conditions.

Constrained optimality conditions. To generate these fixed shares using general portfolio constraints, assume that household i faces some constraint function:

$$G_{it}(\mathbf{a}_{i,t+1}, \mathbf{a}_{it}) \leq 0.$$

This can be interpreted as financial knowledge, participation costs, regulatory requirements, or other microfoundations in the literature. The household's optimality conditions under this con-

straint are:

$$\begin{aligned}
\lambda_{it} &= \frac{u'(c_{it})}{P_t}, \qquad v'(n_{it}) = \lambda_{it} W_{it}, \\
\lambda_{it} Q_t &= \beta_i \nabla_a V_{i,t+1} - \left[\frac{\partial G_{it}(\mathbf{a}_{i,t+1}, \mathbf{a}_{it})}{\partial \mathbf{a}_{i,t+1}}\right]^\top \mu_{it}, \qquad t < T, \\
\nabla_a V_{it} &= \lambda_{it}(Q_t + D_t) - \left[\frac{\partial G_{it}(\mathbf{a}_{i,t+1}, \mathbf{a}_{it})}{\partial \mathbf{a}_{it}}\right]^\top \mu_{it}, \\
G_{it} &\leq 0, \qquad \mu_{it} \geq 0, \qquad \mu_{it}' G_{it} = 0.
\end{aligned} \tag{A.4}$$

The terminal condition is $\mathbf{a}_{i,T+1} = 0$. Substituting the next-period envelope condition yields the asset pricing equation

$$\begin{aligned}
Q_t = M_{i,t+1}(Q_{t+1} + D_{t+1}) & \\
- \frac{1}{\lambda_{it}} \left[\frac{\partial G_{it}(\mathbf{a}_{i,t+1}, \mathbf{a}_{it})}{\partial \mathbf{a}_{i,t+1}}\right]^\top \mu_{it} & \\
- \frac{\beta_i}{\lambda_{it}} \left[\frac{\partial G_{i,t+1}(\mathbf{a}_{i,t+2}, \mathbf{a}_{i,t+1})}{\partial \mathbf{a}_{i,t+1}}\right]^\top \mu_{i,t+1}, &
\end{aligned} \tag{A.5}$$

where $M_{i,t+1} = \beta_i \lambda_{i,t+1}/\lambda_{it}$ is the nominal discount factor.

Proof. For real expenditure $s > 0$, purchasing $P_t s/(Q_t'\omega)$ units of a bundle delivers next-period liquid resources $sR_{i,t+1}(\omega, \epsilon)$. Since continuation value is increasing in liquid resources, portfolio choice maximizes this return independently of s. Both problems therefore reduce to the same scalar policy, evaluated at $R^*(\epsilon)$ and $R^{\text{fix}}(\epsilon) = R(\bar{\omega}_i, \epsilon)$.

At each date, the optimized return weakly exceeds the fixed-bundle return, with equality at $\epsilon = 0$. Their difference therefore has a local minimum at zero. Differentiability in a two-sided neighborhood implies

$$\frac{d}{d\epsilon}\left[R^*_{i,t+1}(\epsilon) - R^{\text{fix}}_{i,t+1}(\epsilon)\right]\Big|_{\epsilon=0} = 0.$$

The finite vectors of return inputs thus differ by $o(|\epsilon|)$, while all other inputs coincide. Differentiability of the reduced policy gives

$$c_{i0}^{\text{endog}}(\epsilon) - c_{i0}^{\text{fixed}}(\epsilon) = o(|\epsilon|).$$

Continuity and positive baseline expenditure keep the positive-scale restriction inactive locally. □

Neither stationarity nor a unique return maximizer is required; date-specific baseline bundles yield equivalence to a fixed schedule of compositions. Negative investment scales, binding quantity or withdrawal limits, and nondifferentiable optimized returns can invalidate the reduction. A utility envelope alone does not establish consumption equivalence.

A.6 Proof of Proposition 2

This section derives the aggregate decomposition in Equation 8. Use fixed baseline household weights and assume $0 < C_0 = \mathbb{E}_i[c_{i0}] < \infty$, integrability of the displayed products, and validity of differentiation under the expectation. Define the following terms for notational simplicity:

$$H_i = \sum_k |MV_{i0}^k|, \qquad L_{it} = MPC_i\, q_{it}(c_{it} - Y_{it}) - c_{i0}\varepsilon_{i0t}^h,$$
$$G_i^Q = P_0^{-1} \sum_k a_{i0}^k\, dQ_0^k, \quad G_i^D = P_0^{-1} \sum_k a_{i0}^k\, dD_0^k.$$

Here, H_i is gross absolute exposure, G_i^Q, G_i^D are asset price and payout gains respectively, and L_{it} is the coefficient on the forward return term. Defining log changes as $\bar{\nu}^Q = d\bar{Q}_0/\bar{Q}_0$, $\bar{\nu}^D = d\bar{D}_0/\bar{Q}_0$, and $\bar{x}_t = d\log\bar{\mathcal{R}}_t$, Proposition 2 defines

$$G_i^Q = H_i(\bar{\nu}^Q + \zeta_{i0}^Q), \qquad G_i^D = H_i(\bar{\nu}^D + \zeta_{i0}^D), \qquad d\log\mathcal{R}_{it} = \bar{x}_t + \zeta_{it}^R.$$

Substituting into (7), taking expectations, and dividing by C_0 proves (8).

Next, I establish invariance to the reference asset. Given the same household problem and the shock, the choice of the reference asset determines how each portfolio response is split between a reference component and a household-specific deviation. For a given household, if a new reference asset assigns an extra dollar of gain to the reference component, then the household's idiosyncratic term falls by one dollar. Both terms are then multiplied by the same MPC, so the changes in each channel's consumption contributions cancel. Summing across households therefore leaves the total response in (8) unchanged. That is, for some perturbation in the reference asset that causes changes in prices, dividends, and returns κ^X for each term X, the decomposition becomes:

$$\frac{dC_0}{C_0} = \frac{1}{C_0}\sum_{t=0}^{T}\mathbb{E}_i\left[MPC_i\frac{dY_{it}}{\mathcal{R}_{it}}\right] - \frac{\mathbb{E}_i[MPC_iA_{i0}]}{C_0}\,d\log P_0 + \frac{\mathbb{E}_i[MPC_iH_i]}{C_0}\left[(\bar{\nu}^Q+\kappa^Q)+(\bar{\nu}^D+\kappa^D)\right]$$
$$+\frac{1}{C_0}\mathbb{E}_i\Big[MPC_iH_i\left[(\zeta_{i0}^Q-\kappa^Q)+(\zeta_{i0}^D-\kappa^D)\right]\Big] + \frac{1}{C_0}\sum_{t=1}^{T}\mathbb{E}_i[L_{it}]\,(\bar{x}_t+\kappa_t^R) + \frac{1}{C_0}\sum_{t=1}^{T}\mathbb{E}_i\left[L_{it}(\zeta_{it}^R-\kappa_t^R)\right].$$

Which is invariant to the specification of κ^X.

B Empirical implementation

B.1 MPCs out of income and wealth

For a consumption function $\mathcal{C}_i$, a change in earnings that generates an income path g_{it} will result in the following amount of spending depending on a household's MPC:

$$b_i = \sum_t MPC_{it}^Y g_{it}, \qquad MPC_{it}^Y = \frac{\partial \mathcal{C}_i}{\partial Y_{it}}. \tag{B.1}$$

In the baseline model where all wealth is equally fungible, $MPC_{it}^Y = MPC_i q_{it}$, this can be written equivalently for spending out of changes in asset prices q_{it}:

$$b_i = MPC_i K_i, \qquad K_i = \sum_t q_{it} g_{it}, \qquad MPC_i = b_i / K_i \quad (K_i \neq 0). \tag{B.2}$$

B.2 Data details regarding the CEX and PSID

Consumption imputation. I impute log annual expenditure using the 1981–2023 CEX, excluding 1982–83. Quarterly spending is multiplied by four, and all measurements are in 2010 dollars. There are 80 regressors in an OLS specification to map consumption from the CEX to the PSID: a constant; log food; age multiplied by indicators for $a < 37$, $37 \leq a \leq 47$, $47 \leq a \leq 56$, and $a \geq 56$; three education and three child indicators (1, 2, 3+); log income and employment; log food consumption interacted with these twelve terms and year indicators (1981 omitted); three region and three spouse-education indicators; household size, race, four price indices, spouse age, and spouse employment.

PSID sample and cleaning The PSID sample contains 31,445 paired observations, pulled from 2,280 original families, and 1,227 unemployment events in 1986, 1996, and biennially from 2003 to 2021. I keep households with heads aged 25–62 and positive earnings, food, and imputed expenditure. I truncate outlier households with absolute log food and earnings changes above four. The unemployment instrument records an unemployment event of either previously employed adult in the household. Outcome, earnings interactions, and instrument interactions are residualized on household-level controls: household-size changes, prior head/spouse employment, state-by-year fixed effects, and all controls used in the MPC estimation procedure. Standard errors are clustered by household.

B.3 Permanent and transitory income comparison

In this section, I compare the MPCs estimated following Patterson (2023) with MPCs estimated using the method from Blundell et al. (2008) to test whether these MPCs are consistently estimated

across households. Following Blundell et al. (2008), let residual annual log income satisfy $y_t = p_t + v_t$, $p_t = p_{t-1} + u_t$, and consumption growth satisfy $\Delta c_t = \phi u_t + \psi v_t + \xi_t$. Within cells, assume constant coefficients; mutually orthogonal, serially uncorrelated u_t, v_t with positive, constant variances; and ξ_t orthogonal to the income changes below. Then

$$\phi = \frac{\mathrm{Cov}(\Delta c_t, \Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})}{\mathrm{Cov}(\Delta y_t, \Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})}, \qquad \psi = \frac{\mathrm{Cov}(\Delta c_t, \Delta y_{t+1})}{\mathrm{Cov}(\Delta y_t, \Delta y_{t+1})}.$$

The income sum equals $u_{t-1} + u_t + u_{t+1} + v_{t+1} - v_{t-2}$, so the ratios reduce to $\phi \,\mathrm{Var}(u_t)/\,\mathrm{Var}(u_t)$ and $-\psi \,\mathrm{Var}(v_t)/[-\,\mathrm{Var}(v_t)]$.

The implementation uses residualized two-year log changes $\widetilde{y}_t$ and $\widetilde{c}_t$, rather than replicating the original minimum-distance estimator. With X_i containing the seven slope terms in Table A.1, define $S_t^P = \widetilde{y}_{t-2} + \widetilde{y}_t + \widetilde{y}_{t+2}$ and $S_t^T = -\widetilde{y}_{t+2}$. For $s \in \{P, T\}$, estimate θ_s from

$$\mathbb{E}[X_i X_i' \widetilde{y}_t S_t^s]\theta_s = \mathbb{E}[X_i \widetilde{c}_t S_t^s].$$

Interpretation requires the corresponding income-process and orthogonality restrictions at the two-year frequency; serial correlation or measurement error can invalidate them.

Figure A.1 shows the comparison between the Patterson and BPP methodologies. The average MPCs from the unemployment IV, BPP permanent income, and BPP transitory income changes are 0.508, 0.298, and 0.171, respectively. The BPP estimates consistently are smaller than the MPC estimates from the unemployment IV, but the heterogeneity across groups is largely the same.

B.4 Portfolio categories and aggregation

Let V_{ik} be the signed real holding of asset k by household i. I let holdings be either positive or negative: Gross assets are $A_i^+ = \sum_{k \in \mathcal{A}^+} V_{ik}$, debt is $D_i^- = -\sum_{k \in \mathcal{A}^-} V_{ik}$, net wealth is $N_i = A_i^+ - D_i^-$, and gross absolute positions are $H_i = A_i^+ + D_i^-$. The household shares plotted in Figure 1 use holdings as a share of gross assets V_{ik}/A_i^+. Note that this is for expositional purposes only, as Proposition 2 uses gross absolute positions H_i as the denominator. Table A.2 shows the full mapping from SCF assets to their empirical counterparts.

B.5 Impulse-response identification and units

For each system of SVARs defined in Section 3.4, define $u_t = B\varepsilon_t$ and $g = \mathrm{Cov}(z_t, u_t)$. The relevance assumption and exclusion restriction imply that g is proportional to the monetary shock. Assuming that the two-year yield entry $g_y \neq 0$, the normalized impulse responses at horizon h are:

$$b^{MP} = -0.25\, g/g_y, \qquad \mathrm{IRF}(h) = \Psi_h b^{MP},$$

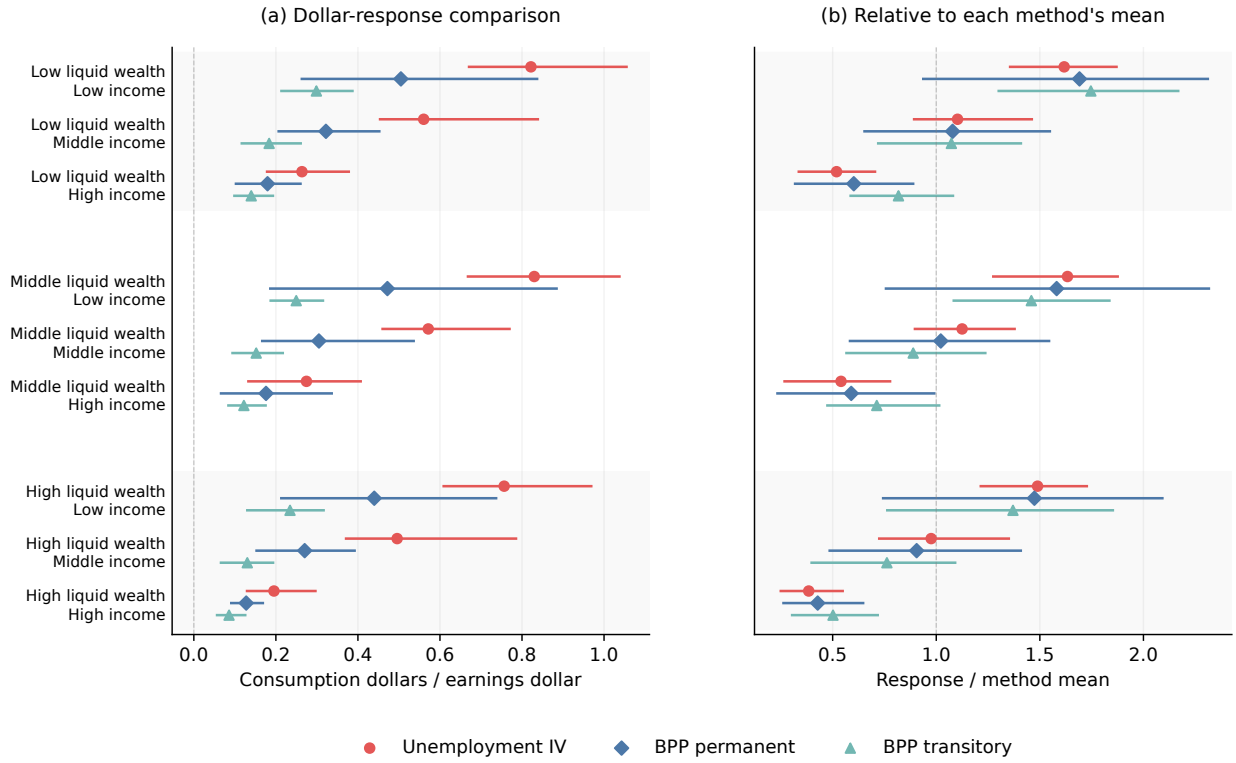


Figure A.1: Unemployment IV and BPP MPC estimates

Notes. Points show consumption responses by income and liquid wealth group, averaging over age. Lines are pointwise 95% confidence intervals. The same PSID sample is used for both estimations. Panel (b) divides each response by its own method's sample mean.

where Ψ_h is the reduced-form moving average coefficient. This follows by taking instrument covariances and rescaling. Responses of $100 \log X$ are divided by 100 to obtain log changes. For some asset k, log differentiation gives

$$d \log R_t^k = \frac{dQ_t^k + dD_t^k}{Q_t^k + D_t^k} - \frac{dQ_{t-1}^k}{Q_{t-1}^k} - (d \log P_t - d \log P_{t-1}).$$

So the cumulative return term is the sum of one-period log changes. The baseline estimation retains separately estimated price, payout and return responses, so this identity is not imposed on all asset classes. Figure A.2 shows selected responses of macroeconomic variables used in the construction of the sufficient statistics. All macro variables used as controls in the SVAR use monthly data from January 1973 through December 2023. The Bauer-Swanson shocks are from 1988-2023.

Payment schedules. I detail some assumptions made on the timing of household asset and debt payments. Let d be the duration on some asset, and let $q = \min\{1, 1/d\}$ be the minimum payment on that asset each year. Thus, principal and coupon payments in year h are $q(1-q)^h$ and $\eta(1-q)^h$, where η is the annual coupon rate. Largely following the duration schedule in Auclert (2019), I set durations to be 0.5 years for CDs and credit-card/other debt; 4 for savings bonds; 7 for government bonds, retirement fixed income, and education debt; 3.5 for corporate bonds; and 3 for vehicle/other installment debt. Government and corporate coupons are 4.5% and 5.5%. Transaction deposits, HELOCs, and other credit lines reset in each month. Taxable equity and other real estate distribute 2% and 4% annually; retirement equity, housing, business, vehicles, and residual assets have no cash payments. Revaluation gains enter once at impact and do not face additional realization adjustments.

Reference choice. The reference asset used in the main empirical results is the wealth-weighted household portfolio. This has the intuition of a representative household's portfolio, as the representative household must hold all assets in the economy. Section A.6 shows that while total monetary policy transmission is invariant to the choice of the reference asset, the decomposition across channels may not be. Figure A.3 compares the main decomposition exercise under the representative portfolio as the reference asset as in the main text, as well as using the zero-coupon Treasury yield from Adrian et al. (2013) or the average portfolio across households as the reference asset. As shown in Appendix A.6, this choice of the reference asset only changes the contribution of each channel to the consumption change and does not affect the total consumption response.

B.6 Comparing to consumption responses in the literature

Because my main estimation strategy flexibly estimates all statistics at the household level before aggregating, it is straightforward to group households into different categories and compare against existing estimates in the literature. First, Figure A.4 shows my main estimation grouped into poor

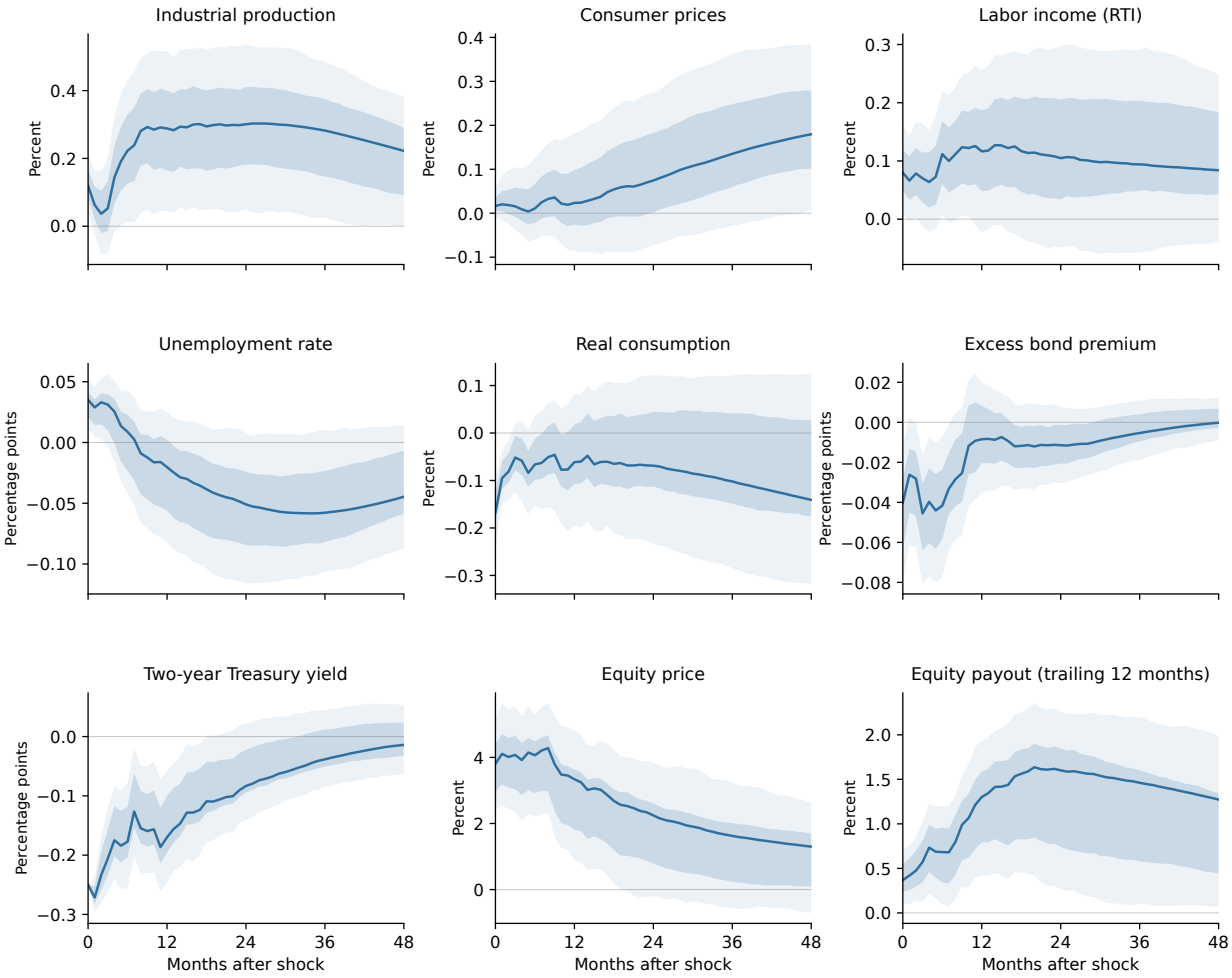


Figure A.2: Additional responses to a 25 basis point decline in the two-year Treasury yield.

Notes. Solid lines are fitted responses; dark/light bands are pointwise 68/95% percentiles of 500 saved wild bootstrap draws. Labor income comes from Realtime Inequality (RTI). Equity payouts are a trailing 12-month dividend proxy.

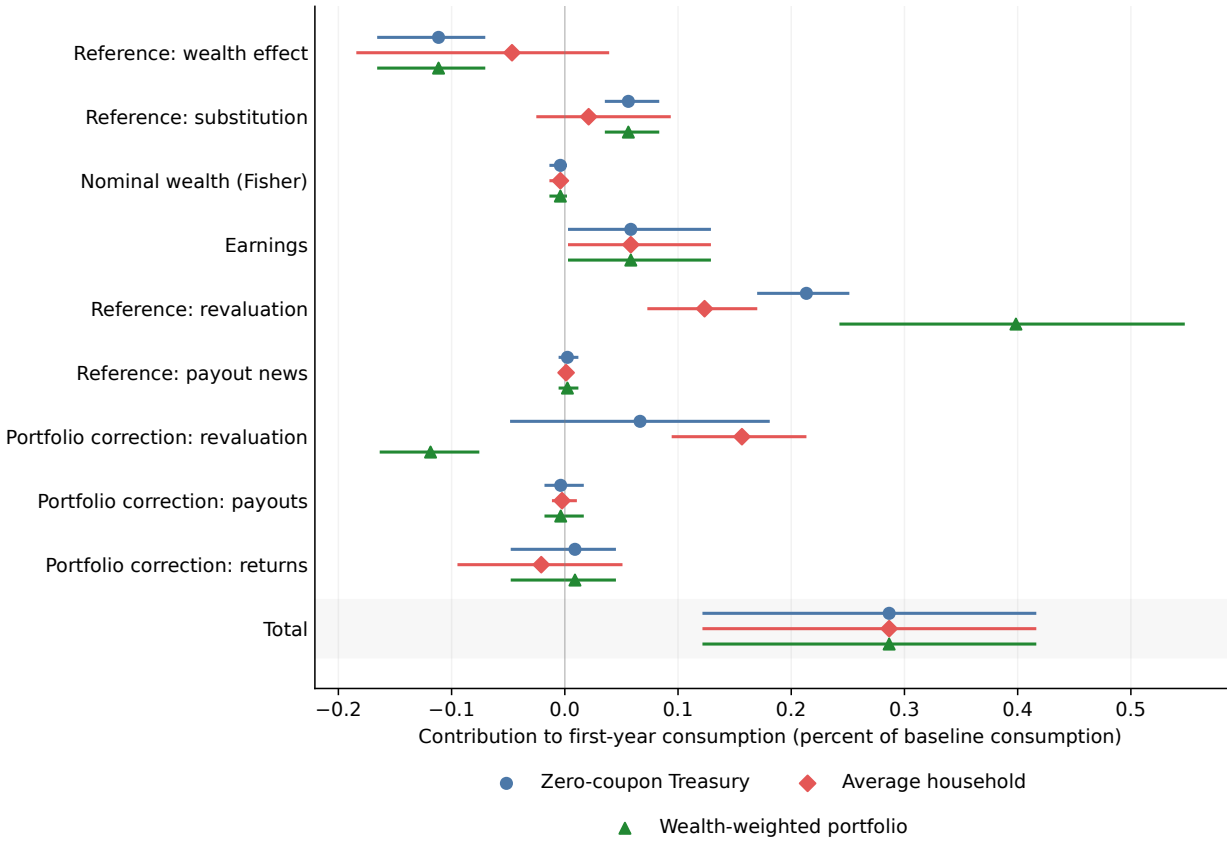


Figure A.3: Sensitivity of channel assignments to the reference asset

hand-to-mouth (HtM), wealthy HtM, and non HtM groups following the categorization in Slacalek et al. (2020). While the top Panel (a) does not exactly match the estimates in Slacalek et al. (2020), as the non-HtM households have a much larger response, the bottom panel using the same MPCs outlined in Slacalek et al. (2020) leads to almost the same distributional response that they find for France in their Figure 7, which is the country that they study with the most similar aggregate impulse responses as the U.S. Thus, the vast majority of the difference between my distributional findings and theirs is the MPCs used in the sufficient statistics. In particular, I use an average MPC of 0.03 from Di Maggio et al. (2020) and consistent with Chodorow-Reich et al. (2021) on capital gains for households in the top 95th-100th percentile of wealth, while they use an MPC of 0.01 on capital gains for non-HtM households.

C Quantitative model and validation

C.1 Households and permanent types

In this quantitative model, I use much of the same notation and structure as in Auclert et al. (2021). Let a, b be current-period illiquid and liquid balances, a', b' be next-period illiquid and liquid balances, and e be an idiosyncratic earnings state with transition matrix Π. Each wealth type j has mass ν_j and discount factor β_j. Union hours N_t are taken as given by households. The household's value function is:

$$V_{jt}(e,b,a) = \max_{c,a',b'} \left\{ u(c) - v(N_t) + \beta_j \sum_{e'} \Pi(e,e') V_{j,t+1}(e',b',a') \right\},$$
$$c + a' + b' + \Phi_j(a',a;r^a_{jt}) = (1-\tau_t) w_t N_t e + (1+r^a_{jt})a + (1+r^b_t)b, \qquad a' \geq 0, \quad b' \geq \underline{b}. \tag{C.1}$$

Adjusting illiquid balances incurs the quadratic adjustment cost:

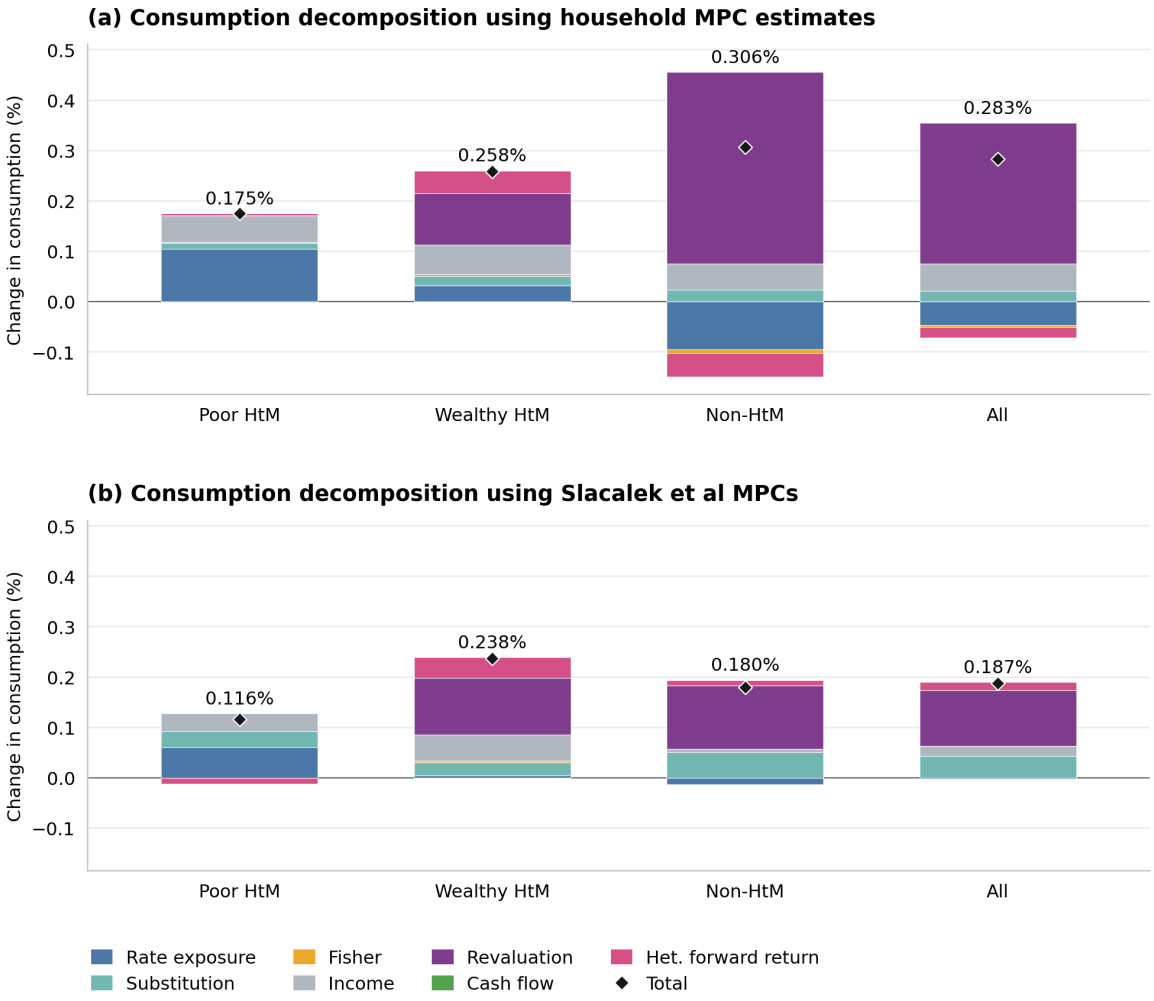
$$\Phi_j(a',a;r^a) = \frac{\chi_{1,j}}{2} \frac{[a' - (1+r^a)a]^2}{(1+r^a)a + \chi_0}, \qquad \chi_{1,j} \geq 0, \quad \chi_0 > 0,$$

Production, capital adjustment, nominal price and wage setting, and monetary policy follow the same implementation as in Auclert et al. (2021). The full calibration is detailed in Table A.3.

C.2 Asset shares and market clearing

Let s_{jt} denote the equity share chosen at t for the payoff at $t+1$, so $\theta_{j,t+1} = s_{jt}$ in (24). Let A_{jt} be the average balance in the illiquid account for type j, $A_t = \sum_j \nu_j A_{jt}$, and B_t be the liquid balance. With the given equity supply p_t, private debt B^p, and government debt B^g, market clearing for assets requires:

$$A_t + B_t = p_t + B^g + B^p, \qquad \sum_j \nu_j s_{jt} A_{jt} = p_t. \tag{C.2}$$



Top: Household estimated MPCs + Di Maggio et al (2020) adjustments.
Wealth-weighted average capital gains MPC (following Di Maggio et al (2020)) = 0.0303.

Bottom: Slacalek et al (2020) MPCs.
income/cash MPCs = 0.50 / 0.50 / 0.05; housing/stock MPCs = 0 / 0.07 / 0.01.

Figure A.4: Comparing the distribution of responses with different MPCs

To satisfy equity clearing, the share of equities for each group j is $s_{jt} = \min\{1, \max\{0, \bar{\theta}_t + \vartheta_j\}\}$. Table A.3 reports the permanent equity types and steady-state shares.

Assume $B_{t-1} \geq 0$. Using the government and firm budget identities and the liquidity transformation cost ω paid by the financial intermediary,

$$\tau_t w_t N_t = G_t + r_t B^g, \qquad d_t = Y_t - w_t N_t - I_t - \Phi_t^p - r_t B^p,$$

summing across household budgets and imposing market clearing gives

$$Y_t = C_t + I_t + G_t + \mathcal{X}_t + \Phi_t^p + \omega B_{t-1},$$

where $\mathcal{X}_t$ is household account-adjustment expenditure and I_t includes capital-adjustment expenditure. Because the equity wedge is caused by the financial intermediary's portfolio constraints, it does not need to be accounted for in market clearing.

For a household's net worth quintile g, let W_g denote total net worth and E_g be broad equity, which includes both public stock exposure and private business wealth from the SCF. I target the wealth shares $w_g = W_g / \sum_h W_h$ as well as the equity shares $e_g = E_g / \sum_h E_h$ for each group. Given the calibrated aggregate ratio of broad equity to household net worth $\bar{s}_E = 0.447$, the implied ratio of equities for each quintile must scale to clear the equity market, so $r_g = \bar{s}_E e_g / w_g$ when $w_g > 0$.

C.3 Incidence relative to own consumption

Figure A.5 retains the same structure as Figure 13, but the right panel divides each quintile's spending response by its *own* baseline consumption instead of as a share of the aggregate change in consumption. This change leads to a more pronounced U-shape: Consumption is relatively higher for the lower wealth quintiles, lower for the middle wealth groups, and higher for the highest wealth group. In this plot, none of the model lines match the data well, and particularly overestimate the response at the bottom. This is for two reasons. The first is that the income channel is substantially larger in all HANK models than in the data, partly because taxes adjust much more quickly than in the data. The second is that MPCs are overestimated in the model, especially in the wealth-types model as lower wealth groups mechanically have lower βs. Additionally, the consumption measure used in the denominator in the data is imputed into the SCF at coarse income groups, which lowers the percent change for those in lower wealth quintiles; ongoing work will address this issue and improve the own-consumption measurement.

C.4 Validation and approximation error

This section details the validation exercise in Section 4. Recall that the quantitative model is quarterly. Define household consumption in the first year following the shock as $\mathcal{C}_i^{(4)} = \mathbb{E}[\sum_{q=0}^{3} c_{iq} \mid i]$. In the empirical decomposition implementation, MPCs are resource-specific as in Section 3:

Table A.1: Household MPC estimation

A. Interacted unemployment-IV slope function

Term	Estimate	95% interval	Wald $\chi^2(7)$	Partial R^2
Constant (low income, low liquidity)	0.765	[0.578, 0.951]	321.9	0.0029
Middle income	-0.255	[-0.440, -0.071]	208.2	0.0031
High income	-0.543	[-0.699, -0.388]	140.9	0.0027
Middle liquid wealth	0.011	[-0.135, 0.158]	146.1	0.0097
High liquid wealth	-0.069	[-0.185, 0.048]	64.1	0.0016
Age minus 45	0.001	[-0.008, 0.009]	161.4	0.0019
(Age minus 45)^2/1,000	0.516	[-0.382, 1.415]	214.8	0.0022

B. Source coefficients (SEs) and income-coefficient ratios

Wealth pct.	Gain	Dividend	Income	Gain/income	Div./income
0–50	0.233 (0.026)	0.560 (0.109)	0.766 (0.004)	0.304	0.731
50–70	0.068 (0.008)	0.601 (0.050)	0.786 (0.012)	0.087	0.765
70–90	0.037 (0.005)	0.512 (0.038)	0.690 (0.020)	0.054	0.742
90–95	0.027 (0.007)	0.587 (0.057)	0.696 (0.004)	0.039	0.843
95–100	0.028 (0.007)	0.386 (0.030)	0.661 (0.014)	0.042	0.584

Notes. Panel A reports coefficients from the MPC estimation in Equation 17 on each relevant observable. Panel B reports the average MPCs that are used for each wealth group, using the Di Maggio et al. (2020) normalization for the capital gain and dividend MPCs.

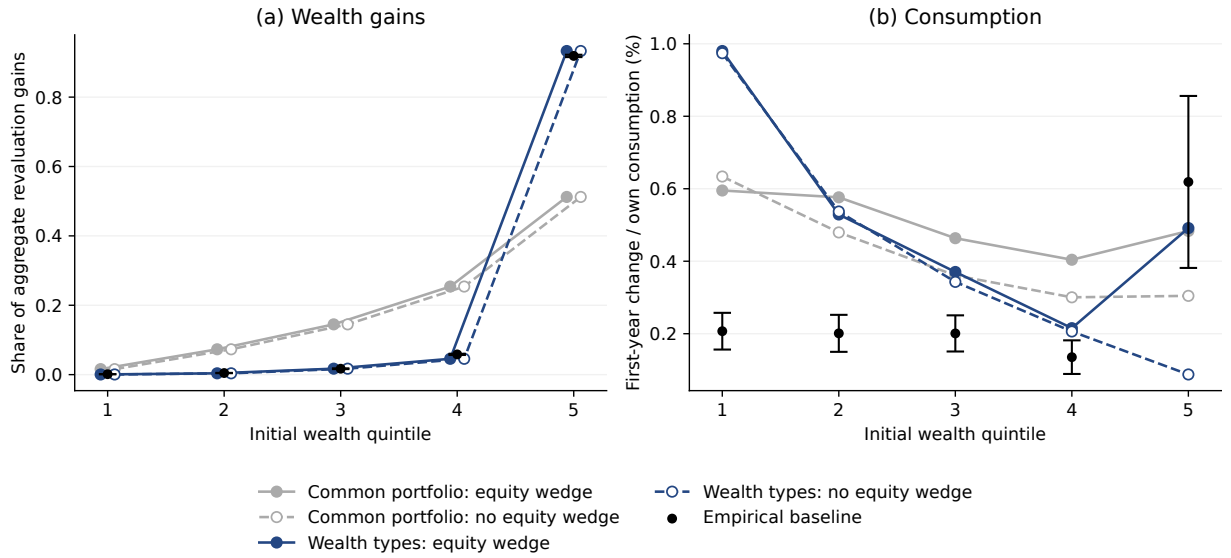


Figure A.5: The counterpart to Figure 13 with group-specific consumption as the denominator at the right panel. Panel (a) retains shares of aggregate revaluation gains. Panel (b) reports the change in first year consumption as a percentage of each quintile's own baseline consumption

Table A.2: Asset categories and empirical mappings

Category	SCF combined-extract fields	Empirical series
Transaction accounts	LIQ (includes money-market accounts)	3-month T-bill return
Certificates of deposit	CDS	National CD rate
Savings bonds	SAVBND	Treasury return
Government bonds/funds	GOVTBND, MORTBND, GBMUTF	ACM zero curve; CRSP Treasury prices and coupons
Corporate/other bonds	NOTXBND, OBND, TFBMUTF, OBMUTF; $\frac{1}{2}$COMUTF	Moody's Aaa yield (DAAA)
Taxable public equity	STOCKS, STMUTF; $\frac{1}{2}$COMUTF	S&P 500 price; CRSP value-weighted ex-dividend return and dividend proxy
Retirement equity	RETEQ	Equity price/return; dividends retained in account
Retirement fixed income	RETQLIQ-RETEQ	Treasury return and payment ladder
Primary residence	HOUSES	Case–Shiller national HPI (CSUSHPISA)
Other real estate	ORESRE, NNRESRE	CRSP/Ziman REIT index 1000811; appreciation and income distributions
Private business	BUS	0.4 times nominal equity-price news; no new payout/return IRF
Vehicles	VEHIC	Vehicle CPI return proxy
Other assets	ANNUIT, TRUSTS, OTHMA, OMUTF, OTHNFIN; residual financial balances	Reference return
Residential debt	NH_MORT, HELOC, RESDBT	Freddie Mac 30-year mortgage rate
Other household debt	CCBAL, OTHLOC, VEH_INST, EDN_INST, OTH_INST, ODEBT	Bank prime rate (DPRIME)

Notes. Holdings are in 2010 dollars. Unspecified bonds use corporate bonds as a proxy; unspecified mutual funds use equity. Payout proxies use trailing 12-month distributions.

Table A.3: Parameters and calibration of the quantitative model

Parameter or calibration choice	Benchmark	Wealth types
A. Households and portfolios		
Permanent types; population mass	1; 1	5; 0.20 each
Household discount factor, β_j	0.975802	Panel D
Union discount factor, β_{macro}	0.975802	0.965000
EIS; Frisch elasticity, $(1/\gamma, 1/\nu)$	$(0.5, 1)$	$(0.5, 1)$
Account costs, (χ_0, χ_1, χ_2)	$(0.25, 6.072583, 2)$	$(0.25, 4.002129, 2)$
Labor disutility, φ	1.526363	2.397188
Liquid spread; borrowing limit, $(\omega, \underline{b})$	$(0.005, -0.25)$	$(0.005, -0.25)$
Earnings persistence; log standard deviation, (ρ_e, σ_e)	$(0.966, 0.92)$	$(0.966, 0.92)$
Broad equity / total net worth, $p/(A+B)$	0.447	0.447
Common equity component, $\bar{\theta}^{ss}$	0.482870	0.482870
B. Production, nominal rigidities, and fiscal policy		
Productivity; capital share, (Z, α)	$(0.714615, 0.195295)$	$(0.714615, 0.195295)$
Depreciation; capital adjustment, (δ, ϵ_I)	$(0.02, 4)$	$(0.02, 4)$
Price and wage markups, (μ_p, μ_w)	$(1.07545, 1.1)$	$(1.07545, 1.1)$
Price and wage Phillips slopes, (κ_p, κ_w)	$(0.1, 0.01)$	$(0.1, 0.01)$
Safe real rate; inflation, (r^{ss}, π^{ss})	$(0.0125, 0)$	$(0.0125, 0)$
Spending / output; labor tax, $(G/Y, \tau^{ss})$	$(0.2, 0.314066)$	$(0.2, 0.314066)$
C. Monetary experiment and numerical solution		
Taylor coefficient, ϕ_π	1.5	1.5
Policy innovation, $(\Delta r_0^\star, \rho_{MP})$	$(-0.0025, 0.765270)$	$(-0.0025, 0.765270)$
Equity wedge, (χ_ζ, ρ_ζ)	$(0.198291, 0.91196)$	$(0.410773, 0.838319)$
Grid sizes, $(n_e, n_a, n_b, n_{\text{aux}})$	$(3, 384, 160, 96)$	$(3, 384, 160, 96)$
Asset bounds, $a \in [0, 4000]$, $b \in [-0.25, 50]$	Same bounds	Same bounds
Policy / distribution tolerances	10^{-10} / 10^{-13}	10^{-10} / 10^{-13}
Transition horizon (quarters)	300	300

D. Permanent-type profile (wealth-type economy)

Type j	1	2	3	4	5
Discount factor, β_j	0.947544	0.947575	0.947576	0.956504	0.985005
Equity tilt, s_j	-0.148970	-0.148970	-0.148970	-0.148970	0.016357
Equity share, θ_j^{ss}	0.333901	0.333901	0.333901	0.333901	0.499227

Notes: Periods are quarters; steady-state output and hours are one. Levels are in units of quarterly output. Earnings use a three-state Rouwenhorst approximation. The Taylor rule is $i_t = r_t^\star + \phi_\pi \pi_t$, $\Delta r_t^\star = -0.0025 \rho_{MP}^t$, and the fiscal rule is $\tau_t w_t N_t = G + r_t B^g$, with fixed $G = 0.2$, $B^g = 2.8$. Private debt $B^p = 4.942$ is fixed by the government debt-to-GDP ratio alongside the equity share of net wealth. The equity pricing wedge is parametrized as $\zeta_t = \rho_\zeta \zeta_{t-1} + \chi_\zeta \Delta r_t^\star$, with $\zeta_{-1} = 0$. Initial distributions are stationary; terminal values use the steady state.

Table A.4: Broad-equity targets, achieved moments, and equilibrium accuracy

Initial net-worth quintile	1 (lowest)	2	3	4	5 (highest)
A. Share of aggregate total net worth (%)					
SCF target	-0.266	1.207	4.286	10.955	83.818
Benchmark	1.629	7.251	15.530	26.843	48.747
Wealth types	-0.181	0.610	3.906	11.634	84.031
B. Share of aggregate broad equity (%)					
SCF target	0.097	0.387	1.364	4.808	93.344
Benchmark	1.587	7.311	14.522	25.313	51.267
Wealth types	0.048	0.379	1.716	4.569	93.287
C. Broad equity / within-quintile total net worth (%)					
SCF observed	—	14.33	14.23	19.61	49.77
Implied at 44.7% aggregate	—	14.33	14.23	19.62	49.78
Benchmark	43.55	45.07	41.80	42.15	47.01
Wealth types	—	27.80	19.64	17.56	49.62

D. Aggregate levels and targeted price responses Moment	Target	Benchmark	Wealth types
Total net worth, $A + B$	14.00	14.00	14.00
Equity value, p	6.26	6.26	6.26
Government debt, B^g	2.80	2.80	2.80
Private debt, B^p	4.94	4.94	4.94
Productive capital, K	5.59	5.59	5.59
Broad equity / total net worth (%)	44.70	44.70	44.70
Nominal equity-price impact (log %)	2.74100	2.74127	2.74099
First equity-price half-life (quarters)	9.99805	9.98765	9.99721
E. Maximum absolute market-clearing residuals			
Asset market: steady state	0	1.76×10^{-12}	3.98×10^{-12}
Goods market: steady state	0	1.80×10^{-11}	7.05×10^{-11}
Asset market: transition	0	6.31×10^{-12}	6.67×10^{-11}
Equity ownership: transition	0	2.07×10^{-11}	4.08×10^{-11}
Goods market: transition	0	4.43×10^{-11}	6.96×10^{-11}

Notes: Targets and the calibration figure use the same pooled 1989–2025 SCF sample. Broad equity is `EQUITY + BUS`; total net worth is `NETWORTH`, including housing. Wave/implicate cells receive equal weight, with survey weights within cells. Each wealth quintile has equal population mass. Panels A and B are fitted directly. Panel C reports observed and implied ratios, not additional fitting restrictions. Ratios are omitted for for groups with negative net worth, although these groups remain in both distributional targets. The model is calibrated quarterly in steady state.

Income, dividends, and illiquid assets use different consumption derivatives computed in steady state.

Let H denote the dates whose news is retained by the household. I compare four measures of the first-year consumption response:

1. $\widehat{C}^{E,H}$: the response from the empirical sufficient-statistics procedure with H periods of news
2. $C^{\text{lin},H}$: the response from the linearized household model with H periods of news
3. $C^{NL,H}$: the response from the nonlinear household with H periods of news
4. $C^{NL,\text{all}}$: the response from the nonlinear household model, with the full news path

Comparing the first two measures isolates errors caused by the linearization and approximation procedures; comparing the last two isolates the effect of the truncated news horizon in the empirical implementation. All four measures are of first-year consumption, and differ only in the solution method or the information households receive. Adding and subtracting the two intermediate responses proves the identity

$$\widehat{C}^{E,H} - C^{NL,\text{all}} = \underbrace{\widehat{C}^{E,H} - C^{\text{lin},H}}_{\text{sufficient-statistics approximation}} + \underbrace{C^{\text{lin},H} - C^{NL,H}}_{\text{household nonlinearity}} + \underbrace{C^{NL,H} - C^{NL,\text{all}}}_{\text{omitted anticipated news}}. \tag{C.3}$$

Table A.5 reports the errors caused by the approximation using the sufficient-statistics approach as well as model nonlinearities with a 25 basis point innovation in the Taylor rule intercept. Note that Figure 11 only shows the results from the empirical procedure $\widehat{C}^{E,H}$ and the full nonlinear solution $C^{NL,\text{all}}$, not all of the economies described above. In general, the sufficient statistics errors, which are the difference between the empirical implementation and the linearized household's solution, are larger than errors from nonlinearities. Overall, the total error of the empirical approximation approach is small and between 3-9% of the total consumption response.

Table A.5: Validation with matched reporting windows and news

Economy	Nonlinear response (%)	Suff.-stats. approx. error (pp)	Nonlinearity error (pp)	Total error (pp)
Benchmark, no wedge	0.3118	0.0218	0.0016	0.0235
Wealth types, no wedge	0.2167	0.0191	0.0002	0.0193
Benchmark, with wedge	0.5259	0.0133	0.0036	0.0169
Wealth types, with wedge	0.4580	0.0123	0.0033	0.0156

Notes: This table shows the sources of error in the empirical sufficient-statistics approach in each economy following the decomposition in Equation (C.3), with a 25 basis point innovation in the Taylor rule intercept. The sufficient statistics approximation error, which is the difference between the decomposition formula and the solution in the linearized household problem, accounts for the majority of the error. Nonlinearies are smaller.