Optimal Policy Rules in HANK[†]

Alisdair McKay
Federal Reserve Bank of Minneapolis

Christian K. Wolf
MIT & NBER

June 17, 2026

Abstract: We characterize optimal policy rules in a business-cycle model with nominal rigidities and heterogeneous households. The policymaker has access to two instruments: the short-term nominal rate, and lump-sum transfer payments. A policymaker with a traditional dual mandate that targets aggregates uses the same policy rule as is optimal in the textbook New Keynesian model. We then consider a policymaker with a distributional objective. The optimal policy rule now contains an additional term, reflecting the effects of the policymaker's instruments on consumption inequality. Since, in our model, monetary policy only has very limited distributional effects, this additional term does not materially change optimal monetary policy. Fiscal stimulus payments, on the other hand, have strongly progressive effects and are thus well-suited to cushion the distributional effects of cyclical fluctuations.

JEL Codes: E52, E62

[†]Email: alisdair.mckay@mpls.frb.org and ckwolf@mit.edu. We thank Sushant Acharya, Adrien Auclert, Anmol Bhandari, Florin Bilbiie, Eduardo Dávila, Serdar Birinci, Jordi Galí, Émilien Gouin-Bonenfant, Ludwig Straub, as well as four anonymous referees for helpful comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. This work was supported by computational resources provided by the BigTex High Performance Computing Group at the Federal Reserve Bank of Dallas.

1 Introduction

Should household inequality affect the conduct of cyclical stabilization policy? Recent years have seen a surge of interest in this question, with several prominent studies arguing that the design of optimal monetary policy is substantially altered by distributional considerations (e.g., Bhandari et al., 2021; Acharya et al., 2023). In principle, household heterogeneity can affect optimal policy design in two separate ways. First, household inequality could alter the transmission from policy instruments to any given target (e.g., to inflation and output). This change in transmission may affect whether or not a policymaker can attain her given targets, and how policy instruments need to be set to do so. Second, household heterogeneity may also alter the targets themselves. For example, policymakers may want to dampen the distributional effects of shocks to the macroeconomy.

We cast optimal policy problems in a heterogeneous-agent New Keynesian (HANK) environment in linear-quadratic form, closely mirroring the canonical representative-agent New Keynesian (RANK) literature. As familiar from this literature, the solution to the linear-quadratic policy problem takes the form of a forecast target criterion (see Giannoni & Woodford, 2002). Building on McKay & Wolf (2023b), we express the target criterion in terms of the causal effects of the policymaker's instruments on the target variables that appear in her objective function. Using this framework, we explore the two ways in which household heterogeneity can in principle affect optimal policy design. On *transmission*, we show that the optimal targeting rule of a "dual mandate" central bank that sets nominal interest rates to stabilize output and inflation is exactly the same as in the textbook three-equation RANK model. Intuitively, this is so because, in our model, household heterogeneity does not affect the relative output-inflation trade-off. On *targets*, we first of all observe that, in our HANK economy, monetary policy has limited distributional effects. It follows that, even if policymakers were to have distributional targets, pursuing them through monetary policy is unattractive—large policy swings would be needed to attain distributional goals, which in turn have significant aggregate efficiency costs. Optimal monetary policy thus again essentially looks as in RANK. In contrast, fiscal policy—in our analysis in the form of uniform stimulus payments—does have considerable distributional effects, making it an effective tool to offset distributional shocks.

Environment. Our analysis is set in a rich business-cycle model with nominal rigidities and household heterogeneity. Households face idiosyncratic income risk and self-insure by borrowing and saving in capital, short-term bonds, and long-term bonds. The policymaker sets short-term nominal interest rates, pays transfers to households, and finances the expenditure through taxation as well as bond issuance. Zooming out, our model environment is consistent with much of the recent HANK literature in that it substantially enriches aggregate spending decisions relative to textbook RANK models, but keeps frictions in price-setting and labor supply relatively simple, yielding an aggregate Phillips curve as in RANK.

Our normative conclusions about optimal policy design are closely tied to the positive question of how changes in policy affect the economy. By the "sufficient statistics" results of McKay & Wolf (2023b), those causal effects can be summarized in terms of impulse responses to exogenous shifts in policy (i.e., to policy shocks). A first step in our model calibration is thus to match the scale of the aggregate effects of interest rate changes to evidence on monetary shock propagation. Turning next to the cross-section of households, we design the model to be consistent with the most salient transmission channels of monetary policy to household balance sheets. Key model ingredients here are the presence of long-duration assets—capturing an important redistribution channel of monetary policy (see Auclert, 2019)—as well as a relatively modest response of aggregate labor income to the monetary shock. Putting everything together, the model generates fairly evenly distributed effects of monetary policy, with consumption responding by broadly similar percentage amounts across the wealth and income distributions. Finally, for fiscal policy, we ensure that our model features an elevated average marginal propensity to consume (MPC), as well as material cross-sectional dispersion in the consumption response to lump-sum transfers.

Optimal policy analysis. We study optimal policy problems cast in linear-quadratic form (e.g., as in Giannoni & Woodford, 2002; Benigno & Woodford, 2012). Linearization of the private-sector relations of our model economy yields the linear constraints. Our quadratic objective is then either an ad-hoc dual-mandate objective (as in standard RANK analyses), or derived from a social welfare function that seeks to insulate individual-level household consumption from aggregate shocks (i.e., our "distributional" objective).

Dual mandate. We begin our analysis by studying the problem of a conventional "dual-mandate" central banker that seeks to close the output gap and to stabilize inflation. In addition to being of practical relevance for real-world central banks, this loss function allows us to transparently explore the effects of HANK changing policy instrument propagation while fixing policymaker targets. The optimal interest rate target criterion for such a dual-mandate central banker turns out to be *exactly the same* as in a standard RANK environment. The key step in the argument is that household heterogeneity only affects the demand side of the economy (i.e., the "IS" curve). In the optimal policy problem, however, this demand block is a slack constraint: the policymaker can pick an output-inflation allocation subject to the model's Phillips curve, and then simply set nominal rates as necessary to generate demand consistent with the desired allocation.[1]

Our theoretical analysis leaves the door open for household heterogeneity to materially affect the *path* of the policy instrument necessary to implement the desired (and unchanged) output and inflation outcomes. However, as we discipline our economy to be consistent with empirical evidence on monetary shock propagation, our model's predictions for interest rates necessarily align with RANK economies disciplined by the same evidence. This broad counterfactual equivalence (in terms of macroeconomic aggregates) of empirically disciplined RANK and HANK economies is studied further in Caravello et al. (2025).

Adding distributional objectives. We next study the optimal policy problem of a policymaker with distributional targets. Specifically, we consider a policymaker with an objective function that penalizes business cycle-induced fluctuations in consumption shares and thus in inequality; in other words, she seeks to provide insurance against business cycles, but does not wish to alter the economy's long-run steady state.[2] A second-order approximation to this objective function delivers a quadratic loss in the output gap and in inflation (echoing the dual mandate case), plus now also in the distribution of consumption across households. The problem of minimizing this quadratic loss subject to the economy's linear constraints yields an optimal implicit targeting rule that now consists of three terms, trading off the policymaker's ability to use her instruments to stabilize her three targets. The third term of this optimal rule—which, importantly, is the sole difference from the conventional dual-mandate rule—is governed by the distributional incidence of the policymaker's instruments.

Our main finding is that, in our environment, distributional concerns do not materially alter optimal monetary policy conduct relative to the dual-mandate benchmark. The intuition for why is best understood through some extreme special cases. If interest rate movements did not at all affect consumption shares (e.g., as in Werning, 2015), then the third term in

[1]The same logic implies that interest rate policy and stimulus check policy are perfect substitutes, as in Wolf (2025). We stress that these arguments rely on our assumption that the supply-side of the model (i.e., the NKPC) is unaffected by household heterogeneity, as in much recent work in the HANK literature. If the output-inflation trade-off changes, then the forecast target criterion would of course also change.

[2]Formally, our analysis relies on optimality of the steady-state distribution of consumption. Similar to Le Grand et al. (2025) and echoing the inverse optimal taxation literature (e.g., as in Heathcote & Tsujiyama, 2021), we ensure this through a particular choice of planner weights. See Section 6.1 for a detailed discussion.

our optimal targeting rule would be exactly zero, and so the optimal rule would collapse to the dual-mandate case; conversely, if monetary policy had large distributional effects, then distributional concerns would likely swamp price and output stability considerations. As stressed above, in our environment, the distributional effects of monetary policy are quite moderate, with all households gaining from a monetary easing, and *vice versa*. It thus follows that monetary policy is rather ill-suited as a tool to deal with the distributional implications of business-cycle shocks, and so our policymaker does not find it optimal to deviate too far from dual-mandate outcomes. These findings contrast with other recent work on optimal monetary policy with heterogeneous households (e.g., Bhandari et al., 2021; Acharya et al., 2023; Dávila & Schaab, 2022) that instead tends to find an important role for distributional considerations. To relate our analysis to those earlier studies we show that, if we were to alter our model to imply larger distributional effects of monetary policy (as in those papers), then we too would find a more important role for distributional concerns.

We finally turn to optimal stimulus check policy. Here we find that interest rate policy and stimulus checks are, for our distributional policymaker, highly complementary tools: fiscal stimulus payments sharply compress consumption inequality, so they can be used to attain distributional objectives, with monetary policy primarily aimed at aggregate stabilization.

Literature. We contribute to the growing literature on optimal policy in business-cycle models with rich heterogeneity (e.g., see Acharya et al., 2023; Bhandari et al., 2021; Le Grand et al., 2025; Dávila & Schaab, 2022). Conceptually, our key contribution is to characterize optimal policy through forecast target criteria that are expressed in terms of easily interpretable and at least partially measurable policy causal effects. The transparency and tractability of our analysis stems in part from our formulation of the policy objective; our policymaker focuses on providing insurance against aggregate shocks, but would not wish to intervene in the absence of such disturbances. The computation of these optimal policy rules leverages sequence-space representations of equilibria (Boppart et al., 2018; Auclert et al., 2021).[3]

Our analysis provides insights on the insurance role of macroeconomic stabilization policy. As we stressed, under our assumptions on policymaker objectives, we can offer a transparent

[3]By the equivalence of perfect-foresight sequence-space and stochastic linear state-space methods, our targeting criterion also applies to the analogous stochastic linear-quadratic optimal control problem. Sequence-space linear-quadratic policy problems have been used in prior work to derive "optimal policy projections" (Svensson, 2005; De Groot et al., 2021; Hebden & Winkler, 2024). A contemporaneous and complementary "HANK" paper that relies on sequence-space methods is Dávila & Schaab (2022). Those authors do not rely on linear-quadratic approximations, thus providing more general results, but without the transparency afforded by our "sufficient statistics" expressions.

characterization of optimal policy—all that matters here are the policy instrument's effects on inequality. Relative to existing work, our analysis thus reveals that normative conclusions about the interaction between optimal monetary policy and inequality depend crucially on the positive question of whether policy changes have strong distributional effects. We believe this insight helps to reconcile some conflicting findings in the literature, with some analyses concluding that that inequality strongly shapes optimal monetary policy design (e.g., Bhandari et al., 2021; Dávila & Schaab, 2022; Smirnov, 2023) and others concluding that it does not (Le Grand et al., 2025, and here).[4]

Finally, our analysis of optimal joint fiscal-monetary policy extends results in Wolf (2025) and Bilbiie et al. (2024). Wolf (2025) considers a model very similar to the one studied here, but the focus is entirely positive, not normative. In a two-agent environment, Bilbiie et al. (2024) argue that monetary and fiscal policy together can stabilize both aggregate activity as well as the consumption shares of the two types. Our work is complementary: we analyze the optimal monetary-fiscal policy mix in an environment with rich heterogeneity.

Outline. Section 2 presents the environment, with the calibration following in Section 3. We then in Section 4 briefly review the connection between optimal targeting rules and the dynamic causal effects of exogenous changes in policy, before finally turning to our main results on optimal policy design in Sections 5 and 6. We conclude in Section 7.

2 Model

We study a relatively standard HANK economy, with two noteworthy features. First, labor supply is intermediated by labor unions, setting wages subject to adjustment frictions. This will allow us to summarize the supply block of our economy through a Phillips curve, focusing our analysis on the demand-side and inequality implications of household heterogeneity, as in much of the recent HANK literature. Second, households invest in a rich menu of assets, allowing the model to capture key features of household portfolios (such as duration and inflation sensitivity) that shape the redistributive effects of policy.

Time is discrete and runs forever, $t = 0, 1, 2, \dots$, and we study linearized perfect-foresight transition paths. As usual, by certainty equivalence, our solutions will be identical to the

[4] Another important theme of the recent optimal policy literature is that inequality among households introduces a new source of time inconsistency and inflation bias (Nuno & Thomas, 2022; Acharya et al., 2023; Dávila & Schaab, 2022; Yang, 2024). This channel is not present in our analysis, as we construct the social welfare function so that the planner does not wish to intervene in the absence of aggregate shocks.

analogous economy with aggregate risk and solved using conventional first-order perturbation techniques with respect to aggregate variables. Throughout this section, boldface denotes time paths (so e.g., $\boldsymbol{x} \equiv (x_0, x_1, x_2, \dots)'$), bars indicate the model's deterministic steady state ($\bar{x}$), and hats denote (log-)deviations from the steady state ($\widehat{x}$).[5]

2.1 Households

The economy is populated by a unit continuum of ex-ante identical households indexed by $i \in [0, 1]$. Household preferences are given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{it}^{1-\gamma} - 1}{1-\gamma} - \nu\left(\ell_{it}\right) \right], \tag{1}$$

where c_{it} is the consumption of household i and ℓ_{it} is its labor supply.

Households face uninsurable risk to their individual incomes. Let ζ_{it} be an idiosyncratic stochastic event that determines the idiosyncratic component of household i's income at date t. The idiosyncratic event ζ_{it} follows a stationary Markov process. Following Werning (2015) and Alves et al. (2020), we assume there is an incidence function Φ that maps aggregate labor income to individual labor income as a function of ζ_{it}. Letting e_{it} be the labor earnings of the household, we have

$$e_{it} = \Phi(\zeta_{it}, m_t, (1-\alpha)y_t),$$

where m_t is a distributional shock that tilts the incidence function towards or away from high-income households, and $(1-\alpha)y_t$ is aggregate labor income. As we describe in greater detail below, labor as a whole receives a share $1-\alpha$ of the total income y_t. The incidence function Φ thus satisfies $\int e_{it} di = (1-\alpha)y_t$ for any value of m_t and y_t, and so the shock m_t only affects the distribution of labor income, but not the total amount. For the quantitative analysis in Section 6, the shock m_t will be our example of an inequality shock—i.e., a shock that affects aggregate demand through redistribution and precautionary savings motives. Next, labor supply is determined by a labor market union (to be described further below), so hours worked ℓ_{it} are taken as given by the household. Total labor income is taxed at some constant proportional rate τ_y. Finally, households receive a time-varying lump-sum transfer $\tau_{x,t} + \tau_{e,t} e_{it}$. Here the first component of the transfer, $\tau_{x,t}$, is the same for all households and will be manipulated as part of the optimal policy problem; we will refer to it as a "fiscal

[5] We use log deviations for $\{y, c, \ell, w, b, a, 1+r, 1+i, q_b, q_k, \eta\}$ and level deviations for $\{\pi, \tau_x, \tau_e, m\}$.

stimulus payment" as it resembles the real-world stimulus checks that have been used in recent recessions in the U.S. The second component, $\tau_{e,t}e_{it}$, is the "endogenous" component, adjusting slowly over time to maintain long-run budget balance. This component of transfers is proportional to the household's productivity.

Households can save and possibly borrow in a variety of assets with different durations and different exposures to surprise inflation. Due to certainty equivalence and no-arbitrage, the returns on all of these assets must be equal at all dates along the equilibrium transition path *except possibly at* $t = 0$, where revaluation effects can lead to heterogeneous realized returns across households. We let r_t denote the (equalized) return between t and $t + 1$, and furthermore let a_{it} denote the net worth of household i at the beginning of period t (inclusive of interest). The overall budget constraint of household i is then

$$\frac{1}{1 + r_t} a_{it+1} + c_{it} = a_{it} + (1 - \tau_y + \tau_{e,t})e_{it} + \tau_{x,t}. \tag{2}$$

All date-0 revaluation effects will be captured by the initial asset position a_{i0}. We will discuss these revaluation effects later. For now we note that—due to those revaluation effects—the distribution of households over initial states (ζ_{i0}, a_{i0}) is endogenous, and we write it as Ψ_0. Finally, we impose a constraint on total household net worth: $a_{it+1} \geq \underline{a}$, where $\underline{a} \leq 0$.

The solution to each individual household i's consumption-savings problem gives a mapping from paths of aggregate income $\boldsymbol{y}$, real returns $\boldsymbol{r}$, transfers $\boldsymbol{\tau}_x$ and $\boldsymbol{\tau}_e$, shocks $\boldsymbol{m}$, and the initial asset position a_{i0} to the path of consumption $\boldsymbol{c}_i$ for a specific realization of idiosyncratic shocks. Aggregating consumption decisions across all households and integrating out idiosyncratic shocks, we thus obtain an aggregate consumption function $\mathcal{C}(\bullet)$:

$$\boldsymbol{c} = \mathcal{C}(\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\tau}_x, \boldsymbol{\tau}_e, \boldsymbol{m}, \Psi_0). \tag{3}$$

2.2 Technology, unions, and firms

We assume that labor union-intermediated labor supply is subject to nominal rigidities, while the prices of final goods are flexible. We make these particular assumptions on nominal rigidities for two main reasons. First, they allow us to mute the response of the labor share to changes in monetary policy, consistent with empirical evidence on monetary policy shock propagation (see the upcoming discussion in Section 3). Second, they turn out to be sufficient to ensure that the aggregate supply side of our economy simplifies to a textbook New Keynesian Phillips curve (NKPC), facilitating comparisons with existing work in representative-agent economies.

Final goods production. We denote the total amount of final goods produced by y_t, and their price by p_t. Output is produced according to a labor-only production function,

$$y_t = zL_t,$$

where z is a (constant) productivity term, and L_t is a bundle of differentiated labor that can be purchased at nominal wage index w_t^n. Following Auclert et al. (2025), final goods producers have market power, allowing them to charge a markup over marginal cost. It follows that $p_t = w_t^n / [z(1-\alpha)]$, where $1/(1-\alpha) \geq 1$ is the gross markup. Since final goods firms have flexible prices, they are able to keep their markup constant at the desired level. It follows that the real wage is $w_t \equiv w_t^n/p_t = (1-\alpha)z$ and so the labor share is constant at $w_t L_t / y_t = 1 - \alpha$. The remaining share α of output is paid to the owners of the firms, which we interpret as capital income (and so the associated stock of wealth as capital). Finally, we assume that maintaining the firms requires a fixed cost δ; in our calibration we will interpret this fixed cost as steady-state investment and calibrate it accordingly.

Labor supply. A unit continuum of labor market unions intermediate household labor supply. Union j hires ℓ_{jt} units of labor from each household; the common labor supply across all households is then $\ell_t = \int_0^1 \ell_{jt} dj$. Union labor is differentiated, resulting in a composite labor bundle

$$L_t = \left(\int_0^1 \ell_{jt}^{\frac{\eta_t - 1}{\eta_t}} dj \right)^{\frac{\eta_t}{\eta_t - 1}}, \tag{4}$$

where the elasticity of substitution, η_t, is subject to exogenous shocks. The nominal wage of type-j labor is denoted by w_{jt}^n, and the cost-minimizing price index of a unit of the labor aggregate is w_t^n. The demand for type-j labor is

$$\ell_{jt} = \left(\frac{w_{jt}^n}{w_t^n} \right)^{-\eta_t} L_t.$$

Notice that labor market clearing implies

$$\ell_t = \int_0^1 \ell_{jt} dj = \int_0^1 \left(\frac{w_{jt}^n}{w_t^n} \right)^{-\eta_t} dj L_t \equiv d_t L_t,$$

where $d_t \geq 1$ summarizes wage dispersion. When wages are dispersed, households must work d_t hours in order to produce a unit of the labor aggregate.

Unions set their wages subject to nominal adjustment frictions. We assume that there is a constant wage subsidy, correcting both tax distortions as well as the steady-state market power of firms and unions, hence delivering aggregate efficiency of the steady state, as usual. Given separable preferences and with all households supplying an equal number of hours, it follows that they all share a common marginal disutility of hours worked. The marginal utility of consumption, however, is generally not equalized across households. To nevertheless ensure that the supply side of our economy stays close to the familiar representative-agent textbook, we assume that the unions evaluate the benefits of higher after-tax income using the marginal utility of average consumption ($c_t^{-\gamma}$) rather than a weighted average of marginal utilities ($\int_0^1 e_{it} c_{it}^{-\gamma} di$), as for example also previously done in Hagedorn et al. (2019) and Auclert et al. (2021). Given this objective, each union sets its wage in standard Calvo fashion, with probability $1 - \theta$ of updating the wage each period. We show in Appendix A.1 that, under our assumptions, the union wage-setting problem gives rise to the following standard linearized NKPC:

$$\widehat{\pi}_t = \kappa \widehat{y}_t + \beta \widehat{\pi}_{t+1} + \psi \widehat{\eta}_t, \tag{5}$$

where κ and ψ are functions of model primitives, defined in Appendix A.1.

2.3 Asset structure

There are three different assets: a short-term nominal bond, a long-term nominal bond, and capital. As discussed above, by no-arbitrage, all assets will provide the same returns along any equilibrium transition path *except possibly at* $t = 0$. This section presents the date-$t \geq 1$ no-arbitrage relations as well as the date-0 revaluation effects.

Asset returns. The three assets pay out the following real returns.

1. *Short-term bond.* The short-term nominal interest rate is denoted i_t. If p_t dollars are invested in the short-term bond at date t, then the payoff is valued at $(1 + i_t) p_t / p_{t+1} = (1 + i_t)/(1 + \pi_{t+1})$ units of goods in the next period.

2. *Long-term bond.* Households can purchase a unit of the long-term bond for a real price of q_t^b. At time $t + 1$, the household receives a real coupon of $(\bar{r} + \sigma_b)(1 + \pi_{t+1})^{-1}$ and furthermore retains a fraction $(1 - \sigma_b)(1 + \pi_{t+1})^{-1}$ of the asset position, now valued at

$(1-\sigma_b)(1+\pi_{t+1})^{-1}q^b_{t+1}$. The parameter σ_b controls the maturity of the asset, with coupons decaying at rate σ_b, while the coupon scaling factor $(\bar{r}+\sigma_b)$ normalizes the steady-state price of the bond to one. The inflation term captures the fact that inflation reduces the real value of all future nominal payouts.

3. *Capital.* A unit share of capital is traded and can be purchased for q^k_t units of the final good. It follows from our discussion in Section 2.2 that the real payoff of a unit of capital is $\alpha y_{t+1} - \delta + q^k_{t+1}$.

Arbitrage relations. By no-arbitrage, all assets must yield the same expected return at all dates $t \geq 1$. Letting r_t denote the common return on these assets for $t \geq 1$, it then follows from the above discussion that in equilibrium we must have

$$1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}} \tag{6}$$

$$1 + r_t = \frac{(\bar{r} + \sigma_b)(1 + \pi_{t+1})^{-1} + (1 - \sigma_b)(1 + \pi_{t+1})^{-1} q^b_{t+1}}{q^b_t} \tag{7}$$

$$1 + r_t = \frac{\alpha y_{t+1} - \delta + q^k_{t+1}}{q^k_t}. \tag{8}$$

Revaluation effects. At time-0, returns are not necessarily equalized, reflecting the arrival of surprise shocks. The expressions on the right-hand sides of (6), (7), and (8) then give us the realized returns at date-0 with i_{-1}, q^b_{-1}, and q^k_{-1} at their steady state values.

At the end of the pre-period $t = -1$, the households in our economy had positions in short-term bonds, long-term bonds, and equities—positions that are then revalued as the date-0 news arrives. We take this initial distribution of portfolios as given; in fact, for our later quantitative analysis, we will match it directly to empirical evidence on household portfolios. It follows from the above discussion of revaluation effects that the date-0 distribution of household asset holdings inclusive of returns, Ψ_0, depends on π_0, q^b_0, q^k_0, and y_0. We will write this mapping as

$$\Psi_0 = \mathcal{H}(y_0, \pi_0, q^b_0, q^k_0). \tag{9}$$

Note that we can use (9) to substitute out for Ψ_0 in the consumption function (3). Linearizing around the deterministic steady state yields the aggregate consumption function

$$\widehat{\boldsymbol{c}} = \mathcal{C}_y \widehat{\boldsymbol{y}} + \mathcal{C}_r \widehat{\boldsymbol{r}} + \mathcal{C}_x \widehat{\boldsymbol{\tau}}_x + \mathcal{C}_e \widehat{\boldsymbol{\tau}}_e + \mathcal{C}_m \widehat{\boldsymbol{m}} + \mathcal{C}_\Psi \left(\mathcal{H}_y \widehat{y}_0 + \mathcal{H}_\pi \widehat{\pi}_0 + \mathcal{H}_b \widehat{q}^b_0 + \mathcal{H}_k \widehat{q}^k_0 \right), \tag{10}$$

where all of the derivative matrices $\mathcal{C}_\bullet$ and $\mathcal{H}_\bullet$ are evaluated at the economy's deterministic steady state.

2.4 Government

The final actor in our model is the government. The government collects tax revenue, pays out lump-sum transfers, sets the nominal interest rate on the short-term bond, and issues short- as well as long-term bonds. Letting a_t^g denote the value of claims on the government entering period t (inclusive of returns), the government budget constraint becomes

$$\frac{a_{t+1}^g}{1+r_t} = a_t^g + \tau_{x,t} - (\tau_y - \tau_{e,t})(1-\alpha)y_t. \tag{11}$$

Government debt maturity. We assume that the government issues both long-term as well as short-term bonds. For all dates $t \geq 1$ along the perfect foresight transition path, the returns on the two bonds are equalized and so the maturity structure of government debt is irrelevant. At date 0, on the other hand, the maturity structure matters via revaluation, perfectly analogously to our earlier discussion of valuation effects in household portfolios. In particular, the revaluation effects will again be embedded in the date-0 portfolio value a_0^g. The aggregate bond holdings of the household sector are the liabilities of the government, and so the revaluation of the government debt is necessarily the mirror image of the revaluation of the household bond positions.

Policy instruments. We consider the nominal rate of interest i_t and the exogenous component of transfers $\tau_{x,t}$ (i.e., "stimulus checks") as the independent policy instruments of the government, used for business-cycle stabilization policy. We assume that the endogenous component of transfers $\tau_{e,t}$ adjusts gradually to ensure long-term budget balance,

$$-\tau_{e,t} = (\bar{r} + \sigma_\tau)\frac{a_t^g - \bar{a}^g}{(1-\alpha)y_t}. \tag{12}$$

When the stock of bonds outstanding exceeds its steady state level, taxes (negative transfers) are raised to pay interest as well as a portion σ_τ of the outstanding bonds. Given this fiscal feedback rule for taxes, government debt then evolves according to (11).

2.5 Equilibrium

We can now define a linearized perfect-foresight transition equilibrium in this economy.[6]

Definition 1. *Given paths of exogenous shocks* $\{m_t, \eta_t\}_{t=0}^{\infty}$*, a linearized perfect foresight equilibrium is a set of government policies* $\{i_t, \tau_{x,t}, \tau_{e,t}, a_t^g\}_{t=0}^{\infty}$ *and a set of aggregates* $\{c_t, y_t, a_t, \pi_t, r_t, q_t^b, q_t^k, \ell_t\}_{t=0}^{\infty}$ *such that:*

1. *The path of aggregate consumption* $\{c_t\}_{t=0}^{\infty}$ *is consistent with the linearized aggregate consumption function* (10)*, and the path of household asset holdings* $\{a_t\}_{t=1}^{\infty}$ *is consistent with the budget constraint* (2)*, aggregated across households.* a_0 *is determined by the existing portfolios of households, valued at date-0 asset prices as represented by* (9)*.*
2. *The paths* $\{\pi_t, y_t, \eta_t\}_{t=0}^{\infty}$ *are consistent with the Phillips curve* (5)*.*
3. *The paths of* $\{\ell_t, y_t\}_{t=0}^{\infty}$ *satisfy the aggregate production function* $y_t = z\ell_t$*.*[7]
4. *The asset prices* $\{r_t, q_t^b, q_t^k\}_{t=0}^{\infty}$ *satisfy* (6)*,* (7)*, and* (8)*.*
5. *The evolution of government debt* a_t^g *and the endogenous component of transfers* $\tau_{e,t}$ *are consistent with the budget constraint* (11) *and fiscal rule* (12)*.*
6. *The output and asset markets clear, so* $y_t = c_t + \delta$ *and* $(a_{t+1} - a_{t+1}^g)/(1 + r_t) = q_t^k$*.*

In this definition, the policy instrument paths $\{i_t, \tau_{x,t}\}$ are taken as given. Sections 5 and 6 will describe optimal policy problems that determine these instruments.

Equilibrium characterization. We can reduce Definition 1 to a small number of linear relations. Lemma 1 provides this more compact characterization of equilibrium dynamics.

Lemma 1. *Given paths of shocks* $\{m_t, \eta_t\}_{t=0}^{\infty}$ *and government policy instruments* $\{i_t, \tau_{x,t}\}_{t=0}^{\infty}$*, paths of aggregate output and inflation* $\{y_t, \pi_t\}_{t=0}^{\infty}$ *are part of a linearized equilibrium if and only if*

$$\widehat{\boldsymbol{\pi}} = \kappa\widehat{\boldsymbol{y}} + \beta\widehat{\boldsymbol{\pi}}_{+1} + \psi\widehat{\boldsymbol{\eta}}, \tag{13}$$

[6] All statements in Definition 1 thus refer to the linearized versions of the relevant model equations.

[7] Note that we drop the efficiency loss term d_t since it is of second order, and thus does not affect a first-order approximation of the production function around a zero inflation steady state (see Galí, 2015). Price dispersion will, however, affect the policy objective functions in Sections 5 and 6.

$$\widehat{\boldsymbol{y}} \;=\; \tilde{\mathcal{C}}_y \widehat{\boldsymbol{y}} + \tilde{\mathcal{C}}_\pi \widehat{\boldsymbol{\pi}} + \tilde{\mathcal{C}}_i \widehat{\boldsymbol{i}} + \tilde{\mathcal{C}}_x \widehat{\boldsymbol{\tau}}_x + \mathcal{C}_m \widehat{\boldsymbol{m}}, \tag{14}$$

where the linear maps $\{\tilde{\mathcal{C}}_y, \tilde{\mathcal{C}}_\pi, \tilde{\mathcal{C}}_i, \tilde{\mathcal{C}}_x\}$ *are defined in Appendix D.1, and* $\widehat{\boldsymbol{\pi}}_{+1} = (\pi_1, \pi_2, \dots)$.

Lemma 1 reduces the complexity of the equilibrium definition in Definition 1 to two equations: the Phillips curve (13) (which is simply a stacked perfect-foresight version of the original relation (5)); and the IS curve (14), which differs from the consumption function (10) in that it imposes output market-clearing, asset revaluation effects, and feedback effects through the government budget to the endogenous component of transfers, $\tau_{e,t}$. Together, these two equations fully characterize the evolution of output and inflation given exogenous non-policy shocks $\{m_t, \eta_t\}_{t=0}^{\infty}$ and policy choices $\{i_t, \tau_{x,t}\}_{t=0}^{\infty}$.

Discussion. How does our model differ from the canonical representative agent New Keynesian models (Galí, 2015; Woodford, 2003), and how does it relate to recent contributions in the HANK literature? *Positively*, the main change is that a simple aggregate Euler equation,

$$\widehat{y}_t = -\frac{1}{\tilde{\gamma}}(\widehat{i}_t - \widehat{\pi}_{t+1}) + \widehat{y}_{t+1}, \tag{15}$$

is now replaced by the more general aggregate demand block (14).[8] Inequality thus affects the aggregate dynamics of our economy in response to shocks and policy actions only through the demand side, with supply—in particular the Phillips curve (13)—kept as in standard representative-agent models. To arrive at this clean separation, our assumptions on union-intermediated labor supply (see Section 2.2) are central. We adopt this approach because the demand-side effects of household heterogeneity are the focus of the recent HANK literature (e.g., see Kaplan et al., 2018; Auclert et al., 2024).[9] *Normatively*, household inequality may affect social welfare functions and thus change policymaker objectives. The remainder of the paper studies the implications of these two changes for optimal policy design.

[8] Equation (15) combines the standard consumption Euler equation with EIS $1/\gamma$ with the log-linearized aggregate resource constraint $\hat{c}_t = (\bar{c}/\bar{y})\hat{y}_t$. The parameter $\tilde{\gamma}$ that appears in (15) is then $\gamma\bar{c}/\bar{y}$.

[9] As an added benefit, our union structure with uniform labor rationing allows us to sidestep counterfactual cross-sectional labor supply responses to policy interventions (see Auclert et al., 2023).

3 Model calibration and the causal effects of policy

In our subsequent analysis, the causal effects of monetary and fiscal policy on aggregate and distributional outcomes will emerge as formal "sufficient statistics" for optimal stabilization policy design (following McKay & Wolf, 2023b). Our model calibration, presented in this section, is designed to align our model with the relevant empirical evidence on the effects of policy and the strength of key policy transmission channels.

Our model calibration is informed by a number of direct and indirect targets. On the aggregate side, the literature has produced robust evidence on the sign and size of monetary policy causal effects on output and inflation, and we discipline our model to be consistent with that evidence. For fiscal policy, we similarly ensure that our model is consistent with the existing evidence on household marginal propensities to consume (MPC). In the cross-section, some prior empirical work has tried to directly estimate the causal effects of monetary policy on consumption inequality. Results so far, however, are not conclusive, with the lack of consensus likely to reflect the lack of high-quality, high-frequency consumption data.[10] We thus instead adopt an indirect approach, calibrating the model to be consistent with empirical evidence on how monetary policy affects the various components of the household budget constraint—something for which there is much more consensus, as we describe below. We then rely on the structure of our model's consumption-savings problem to map these effects on income and wealth into effects on consumption.

3.1 Model calibration

We begin with a discussion of our model calibration—both calibration targets as well as the resulting parameter values. A more detailed discussion of our model's implications for policy causal effects follows afterwards. The full list of all parameters is provided in Table 3.

Income process. Since many of the distributional effects of monetary policy involve changes in asset values (Andersen et al., 2023), it is important for our model to generate a concentrated distribution of wealth, mirroring the data. It is of course well-known that standard incomplete-markets models of the consumption-savings decision struggle to

[10] Among papers that attempt direct measurement of consumption responses, Holm et al. (2021) find that expansionary monetary policy has moderately pronounced U-shaped effects on consumption in the wealth cross-section (in Norway), Coibion et al. (2017) report declines in consumption inequality, and Chang & Schorfheide (2022) conclude that inequality actually even increases somewhat (both in the U.S.).

generate sufficiently concentrated wealth holdings in the top tail of the distribution, and so we follow Castaneda et al. (2003), Boar & Midrigan (2022), and Greenwald et al. (2021) in specifying an income process with superstar earners, generating a realistic concentration of wealth. A second empirical regularity central for our purposes is that households with less education and low past earnings tend to be more exposed to cyclical fluctuations in labor market conditions (Okun, 1973; Hoynes, 2000; Guvenen et al., 2014; Patterson, 2023). In addition, Amberg et al. (2022) and Broer et al. (2022) present evidence specific to monetary policy shocks, showing that the labor market outcomes of poorer households tend to be more sensitive to monetary policy. We thus want our income process to incorporate such unequal incidence of individual earnings to the changes in aggregate earnings.

With those objectives in mind, we assume that the underlying household income state, ζ_{it}, follows a two-component process, with households being either regular workers or high earners. The function Φ that maps ζ_{it} to labor productivity e_{it} is parameterized as

$$\log e_{it} = \log\left[(1-\alpha)y_t\right] + \log\left(\zeta_{it}\right)\left(1 + m_t + \chi \log y_t\right) + \log \bar{e}_t, \tag{16}$$

where χ controls the sensitivity of income dispersion to the aggregate cycle (with a negative χ implying that low-ζ households are more exposed to the cycle), and $\bar{e}_t$ is a normalization constant ensuring that $\int e_{it} di = (1-\alpha)y_t$ at all dates. For regular workers, the state $\log \zeta_{it}$ follows an AR(1) process. High-earning households instead can receive one of two (high) levels of earnings. There are six parameters associated with the high-earnings states: the probability of becoming a high earner of either kind, the earnings level associated with each state, and the probabilities of transitioning back to being a regular worker.[11]

To estimate the full income process we extend the estimation procedure of Guvenen et al. (2022). Those authors propose a rich parametric income process that allows for differential exposure to aggregate business-cycle conditions, and then estimate it using a variety of moments taken from Social Security Administration data. This strategy pins down χ in (16) using data on the differential impact of business-cycle episodes across the income distribution. We extend their procedure by also targeting moments of the wealth and income distributions, as reported in Table 1. Further details on the calibration of the income process and its parameters are presented in Appendix B.1.

[11] To reduce the number of parameters, we assume that there are no transitions between the high-earning states, and that upon exit households draw from the unconditional distribution of regular ζ's. We further assume that regular workers share a common value of χ, while high earners have $\chi = 0$.

	Wealth		Income	
	Data	Model	Data	Model
Top 1%	37	27	17	20
Top 5%	65	66	32	32
Top 10%	76	82	43	44
Top 25%	91	96	64	60
Top 50%	99	101	84	77

Table 1: Shares (%) of wealth and income concentrated in the top x% of the distribution. Data are from the 2019 Survey of Consumer Finance.

Asset structure. Households have positions in capital, short-term bonds, and long-term bonds. We need to calibrate the total supply of each asset class and then how these assets are distributed across households in steady state. As discussed above, those steady state portfolios will determine the distribution of capital gains at $t = 0$, following a shock.

We begin by using data on the balance sheets of U.S. households from the Financial Accounts and the 2019 SCF to impute holdings of particular asset classes across the household distribution of net worth. The top panel of Table 2 reports data on asset holdings in coarse net worth bins, taken from the Distributional Financial Accounts (DFAs).[12] An important point to note is that pension assets represent 22% of total household assets. Of these pension assets, 61% are defined benefit entitlements. Given that we do not explicitly model pensions we must make a choice about how to incorporate these assets into our calibration. As households are quite insulated from the returns on these pension assets, we choose to incorporate the pension funds on to the government balance sheet. The returns on these pension assets are then gradually paid out to households through adjustments in taxes.

The next step is to map the observed asset categories into the corresponding asset classes in the model. We treat corporate and non-corporate equity as levered claims on capital, so each dollar of equity represents \$1.32 dollars of capital and a \$0.32 debt position. This leverage ratio reflects the ratio of debt to assets in the non-financial business sector. Based on the results in Greenwood et al. (2010) we assume that 62% of this debt is long-term debt. We set the duration of the long-term bond in our model to 10 years, in line with results for corporate bonds in Guedes & Opler (1996). Next, we treat mortgages as a 50-50 split of short-term and long-term debt. Even though many mortgages are 30-year loans, the duration

[12] Our imputation procedure yields similar results to the DFAs, but we use the SCF to create portfolios that are smooth functions of net worth, as opposed to the coarse bins shown in Table 2.

Line	Category	Total	Holdings by net worth group			
			Top 1%	Next 9%	Next 40%	Bottom 50%
1	Real estate and durables	167	24	48	72	23
2	Equity and mutual funds	191	101	66	23	2
3	Currency, deposits, and similar	60	16	23	19	2
4	Govt. and corp. bonds and similar	29	10	11	7	1
5	Pension assets	131	6	63	58	4
6	Mortgage liabilities	49	2	12	24	11
7	Consumer credit and loans	24	1	2	8	12
8	Net worth excluding pension assets	374	147	135	89	4
9	Capital	419	157	135	101	25
10	Short-term bonds	-12	1	7	-3	-16
11	Long-term bonds	-33	-11	-8	-9	-5
12	Total	374	147	135	89	4

Table 2: Assets and liabilities in percent of GDP with decomposition by net-worth. Lines 1-8 report the observed asset and liability categories from the Distributional Financial Accounts. Lines 9-12 map these assets and liabilities into their model counterparts.

of these loans is much shorter due to amortization and prepayment (e.g., Greenwald et al., 2021, report a duration of approximately 5 years). We treat consumer credit as short-term debt. Our choices lead to an overall capital-output ratio of 4.19, see Table 2. Households have debts of 12% of GDP in short-term bonds and 33% of GDP in long-term bonds.[13]

We finally set the capital share to 36%, which requires $\alpha = 0.36$. We also interpret the fixed cost as steady-state investment, which we match to a 17% investment to GDP ratio (the average observed from 1984 to 2019). These values, together with capital-output ratio of 4.19, imply a steady state annual interest rate of 4.6%. We normalize productivity z so that steady state output is 1. Lastly, we set the borrowing limit—which restricts net worth and not gross positions—to -0.27 times steady state average income. This value matches the 5th percentile of net worth in our SCF data.

[13] As households (excluding pension funds) collectively have negative credit positions, the pension funds will lend to the households and to the government. As we consolidate the pension funds and government in our model, it is thus as if the government is lending to the households. See Appendix B.2 for further description of how we use the SCF to impute asset positions as smooth functions of household net worth.

Parameter	Description	Value	Calibration target
Households			
$\{\zeta_{it}\}$	Income risk process	–	See text
χ	Income exposure	-2.0	Het. earnings cyclicality
γ	Relative risk aversion	1.2	Monetary shock effects
ϕ	Frisch elasticity	1	Standard
β	Discount factor	0.984	Asset market clearing
Firms			
κ	Phillips curve slope	0.04	Monetary shock effects
$\bar{\eta}$	Labor Substitutability	6	Basu & Fernald (1997)
Asset Structure			
α	Capital share	0.36	Standard
$\delta/\bar{y}$	Maintenance rate	0.17	Investment rate
$\underline{a}/\bar{y}$	Borrowing limit	-0.27	Fifth percentile of net worth
σ_b	Long-term bond duration	0.025	Guedes & Opler (1996)
Government			
τ_y	Labor tax rate	0.23	Steady-state budget balance
$\frac{\bar{\tau}_x}{\bar{y}}$	Transfer share	0.17	After-tax vs. before-tax income
σ_τ	Tax-debt responsiveness	$= \sigma_b$	See text

Table 3: Calibration of our quantitative HANK model. The model period is one quarter.

Fiscal system. The U.S. fiscal tax-and-transfer system is reasonably well-approximated by a uniform baseline lump-sum transfer coupled with a constant marginal tax rate (see Kaplan et al., 2018). Data on post-government and pre-government income from the Congressional Budget Office (2019) imply a steady state transfer of 0.17 times average income, pinning down the ratio $\bar{\tau}_x/\bar{y}$ in our model.[14] The speed of fiscal adjustment is controlled by σ_τ, which we set equal to σ_b, implying very gradual fiscal feedback.

Other parameters. We set the Frisch elasticity to one, and the elasticity of substitution between labor varieties to six, based on Basu & Fernald (1997). Finally, we set the coefficient

[14]We exclude Medicaid and CHIP benefits from the after-tax income. We then regress after-tax incomes on before-tax incomes for the first four quintiles of the income distribution. The intercept of this regression gives us the steady state transfer.

of relative risk aversion γ and the slope of the Phillips curve κ to match the peak responses of output and inflation to an empirically identified monetary policy shock (following Gertler & Karadi, 2015), ensuring that our model is consistent with overall aggregate causal effects of changes in monetary policy. The targeted empirical evidence of a muted price response as usual calls for a flat NKPC; in our case, the implied value of the Calvo parameter θ is 0.886. As usual, matching not just the level but also the precise shape of identified impulse responses to monetary shocks would require several further model ingredients, including in particular habits in consumption (or inattention) as well as price indexation. We omit those ingredients in the interest of parsimony, transparency, and comparability with standard optimal policy analyses in RANK. The details for the internal calibration are provided in Appendix B.1.

3.2 The causal effects of policy

We now give further details on our calibrated model's implications for the causal effects of policy changes on aggregate and distributional outcomes. Those causal effects will loom large in our subsequent optimal policy analysis in Sections 5 and 6.

Aggregates. Impulse responses of real rates, aggregate output, and inflation to monetary policy shocks are displayed in Figure B.2. We generate a policy shock to perfectly match the impulse (i.e., the real rate path) to the empirical evidence of Gertler & Karadi (2015), and our model calibration then ensures that the (peak) aggregate consumption and inflation responses also match the data.

For fiscal policy (i.e., stimulus checks), the direct spending response is simply a function of the economy's average MPC, which in our environment is equal to 17%, broadly consistent with existing empirical evidence (e.g., Parker et al., 2013; Fagereng et al., 2021). In general equilibrium, by demand equivalence (Wolf, 2023), this impulse to demand will be propagated through the same mechanisms as the monetary policy shock discussed above.

Distribution. We next turn to the distributional effects of policy. Figure 1 shows the model-implied impact change in consumption across the wealth distribution following an expansionary monetary policy shock, with the shock to the real interest rate again taken from our estimates based on Gertler & Karadi (2015). Our main finding is that households across the wealth distribution increase their consumption relatively uniformly in percentage terms (total effect line, in dark blue), implying small effects on consumption *shares*.

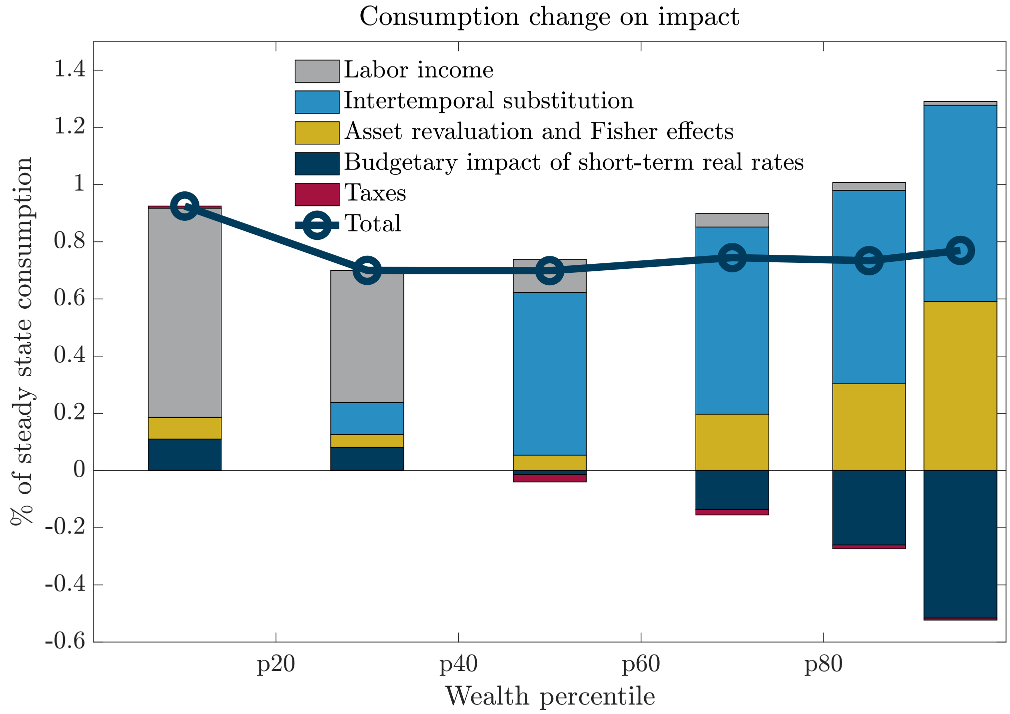

Figure 1: Initial response of consumption to an expansionary monetary policy shock across the distribution of wealth. The policy shock is identical to that estimated in the data (see Appendix B.1). The decomposition shows the effect of each of the aggregate variables that enter the household's decision problem. For real interest rates we split the total effect into "intertemporal substitution" (i.e., the effect of the interest rate purely as a relative price in the Euler equation) and "budgetary impact" (i.e., the effect of the interest rate on the household budget constraint).

What explains these rather evenly distributed effects of monetary policy stimulus? To understand what is going on, the different bars in Figure 1 decompose consumption changes into various forces. On one end of the spectrum, low-wealth households tend to also have rather low incomes, their labor earnings are more exposed to aggregate income, and they are often borrowing-constrained, thus implying high MPCs. As a result, changes in their labor earnings pass through strongly into consumption. These households also tend to have meaningful debts, so low interest rates reduce interest expenses and inflation reduces the real repayment. As we then move towards higher wealth levels, income effects start to matter less and intertemporal substitution effects become stronger. The key thing to note here is that high-wealth households are actually well-hedged with respect to changes in interest rates, as reflected in the asset revaluation effects (yellow) being very similar in magnitude to the negative effect of the lower return on savings (navy blue). Intuitively, wealthy households smooth their consumption, and so their assets fund consumption far into the future. Interest rate changes both affect the value of those long-duration consumption plans (i.e., the navy blue effect) and the values of the similarly long-term assets (the yellow effect), in particular capital. The overall consumption response thus largely reflects substitution effects.[15]

We stress that two particular features of our HANK model play an important role in generating the response profiles of Figure 1. First, the composition of household portfolios is key. Andersen et al. (2023) present empirical evidence that expansionary monetary policy shocks induce substantial capital gains at the top of the income distribution. Gains of that type are captured in the "asset revaluation and Fisher effects" category shown in Figure 1. An alternative model with only short-term nominal bonds would instead have a counterfactually negative contribution from this category, as only the Fisher effect would be present. Second, our sticky-wage, flexible-price model implies a constant labor share. This is consistent with our empirical monetary policy shock evidence, which implies an elasticity of earnings with respect to aggregate output just above one (Figure B.3). Later, in Section 6.4, we contrast our sticky-wage model with a sticky-price variant in which the labor share moves much more. In that alternative model, since labor earnings increase while profits drop, monetary policy is much more progressive than in our baseline environment (see Figures B.3 and B.4).

While monetary policy in our model is thus approximately distributionally neutral, stimulus checks are instead very progressive. The intuition is straightforward: the induced general equilibrium effects of stimulus checks are broadly similar to those of a monetary policy easing (because of demand equivalence), but the incidence of the direct impetus to demand is very different, now being concentrated among the high-MPC poor. A detailed breakdown of the level and composition of the spending response to stimulus checks is provided in Figure B.5. This difference in distributional effects will later play a key role in our conclusions about optimal distributional policy for a policymaker using both instruments.

4 From policy causal effects to optimal policy rules

This section provides a brief review of linear-quadratic optimal policy problems in the sequence space, with particular emphasis on the role of policy causal effects in shaping optimal policy design. The discussion builds heavily on our earlier work in McKay & Wolf (2023b),

[15] In McKay & Wolf (2023a), we construct a similar figure (Figure 3 there) using an empirical approach as opposed to a model-based simulation. There we estimate how various components of income and wealth respond to monetary policy shocks and then make assumptions about the MPCs out of these changes based on existing estimates. Those results also show a broad-based incidence of policy, but with a slightly stronger response among wealthy households.

but is designed to be self-contained. Our subsequent analysis—for both the dual mandate and distributional objectives—will fit into this general framework.[16]

4.1 Linear-quadratic problems in the sequence space

We consider a policymaker with quadratic loss function

$$\mathcal{L} \equiv \frac{1}{2}\sum_{t=0}^{\infty}\beta^t \sum_{i=1}^{I}\lambda_i x_{it}^2 = \frac{1}{2}\boldsymbol{x}'\left(\Lambda \otimes W\right)\boldsymbol{x}. \tag{17}$$

The policymaker targets I variables, and x_{it} is the deviation of the ith target variable from its target value at date t. As before boldface denotes perfect-foresight sequences, β is the policymaker discount factor, and the λ_i's are the policy-target weights. For notational compactness we stack the variables and coefficients as $\boldsymbol{x} \equiv (\boldsymbol{x}_1, \boldsymbol{x}_2, \dots, \boldsymbol{x}_I)'$, $\Lambda \equiv \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_I)$, and $W = \text{diag}(1, \beta, \beta^2, \dots)$, respectively. Our analysis in Sections 5 and 6 will consider particular policymaker objectives that can be written in the quadratic form (17).

The policymaker faces constraints imposed by the equilibrium relationships between variables. These linear constraints are expressed compactly as

$$\mathcal{H}_x\boldsymbol{x} + \mathcal{H}_z\boldsymbol{z} + \mathcal{H}_\varepsilon\boldsymbol{\varepsilon} = \boldsymbol{0}, \tag{18}$$

where $\boldsymbol{z} \equiv (\boldsymbol{z}_1, \boldsymbol{z}_2, \dots, \boldsymbol{z}_J)'$ stacks time paths for the J policy instruments available to the policymaker, and $\boldsymbol{\varepsilon} \equiv (\boldsymbol{\varepsilon}_1, \boldsymbol{\varepsilon}_2, \dots, \boldsymbol{\varepsilon}_Q)'$ similarly stacks the paths for Q exogenous shocks. $\{\mathcal{H}_x, \mathcal{H}_z, \mathcal{H}_\varepsilon\}$ are then conformable linear maps. In our analysis, those constraints are simply the various linear relations that summarize the equilibrium of the economy in Section 2, as stated in Definition 1. Crucially, by McKay & Wolf (2023b), we can equivalently and more conveniently express these constraints as

$$\boldsymbol{x}_i = \sum_{j=1}^{J}\Theta_{x_i,z_j}\boldsymbol{z}_j + \sum_{q=1}^{Q}\Theta_{x_i,\varepsilon_q}\boldsymbol{\varepsilon}_q, \quad i = 1, 2, \dots, I \tag{19}$$

where the Θ's are linear maps that capture the dynamic causal effects of a policy instrument path $\boldsymbol{z}_j$ or shock path $\boldsymbol{\varepsilon}_q$ on a target variable path $\boldsymbol{x}_i$. The alternative constraint (19) thus

[16] As usual, the solutions to the linear-quadratic sequence-space policy problems that we study here are equal to impulse response functions to stochastic shocks under optimal commitment policy in the analogous linearized economy with aggregate risk. See in particular Hebden & Winkler (2024, Section 3.2).

expresses the policy targets directly in terms of impulse responses to policy instruments and exogenous shocks, as opposed to imposing implicit relationships as in (18).[17]

Optimal policy. The optimal policy problem is now to choose the instrument paths $\boldsymbol{z}$ to minimize (17) subject either to (18) (for the original constraint formulation) or (19) (for the re-cast constraint); we will consider the more transparent and economically instructive formulation (19). The resulting first-order condition of this problem is

$$\sum_{i=1}^{I} \lambda_i \cdot \underbrace{\Theta'_{x_i,z_j} W}_{\text{(discounted) causal effect of } z_j \text{ on } x_i} \cdot \boldsymbol{x}_i = \mathbf{0}, \quad j = 1, 2, \ldots, J. \tag{20}$$

(20) says that, for each instrument $\boldsymbol{z}_j$, the paths of the targets $\boldsymbol{x}_i$ must be at an optimum within the space implementable through $\boldsymbol{z}_j$. In the language of Svensson (1997) and Woodford (2003), this rule is an example of a so-called implicit "target policy criterion": the policymaker sets the available instruments to align projections (i.e., future paths) of macro aggregates as well as possible with its targets, given what is achievable through the available instruments. We finally note that the corresponding optimal policy path is then given as

$$\boldsymbol{z}^* \equiv -\left(\Theta'_{x,z}(\Lambda \otimes W)\Theta_{x,z}\right)^{-1} \times \left(\Theta'_{x,z}(\Lambda \otimes W)\Theta_{x,\varepsilon}\boldsymbol{\varepsilon}\right), \tag{21}$$

where $\Theta_{x,z}$ and $\Theta_{x,\varepsilon}$ suitably stack the individual Θ_{x_i,z_j}'s and $\Theta_{x_i,\varepsilon_q}$'s. Intuitively, the policymaker will use her tools as well as possible to offset the disturbances $\boldsymbol{\varepsilon}$, with the optimal policy response taking a simple weighted least-squares regression form.

4.2 Connection to empirical evidence

For our purposes, a key takeaway of the preceding discussion is that optimal policy design in the linear-quadratic setting is fully characterized by two objects: the policymaker preferences, and the causal effects of policy instruments z on policymaker targets x, as expressed in $\Theta_{x,z}$.

In the next two sections we will consider two optimal policy problems, where the targets

[17] The equivalence of (18) and (19) would be immediate for invertible $\mathcal{H}_x$. In typical macroeconomic models, however, $\mathcal{H}_x$ is not invertible, so recasting the constraint as (19) requires additional arguments. McKay & Wolf (2023b) provide those arguments; briefly, the core intuition is that the optimal policy problem can be shown to be equivalent to the alternative, artificial problem of picking shocks to a given baseline, determinacy-inducing policy rule. Policy and non-policy shocks relative to this arbitrary baseline policy rule then yield the impulse response matrices Θ. See Appendix C.1 for further details.

x will be output, inflation, and household consumption shares. In principle, one could hope to leverage the results of this section by directly measuring the causal effects of monetary and fiscal policy on these targets. Doing so in a fully reduced-form way is challenging, however, in particular for the distributional effects of monetary policy, for reasons discussed in Section 3. We will thus instead use our calibrated model to deliver the $\Theta_{x,z}$'s, but note that by design of the calibration that model is consistent with key pieces of salient empirical evidence.

5 Dual mandate

We begin with the problem of a conventional dual-mandate policymaker—i.e., a policymaker that simply seeks to stabilize fluctuations in inflation and in the output gap. Such a loss function is not just of practical relevance, but also conceptually serves as a useful benchmark: heterogeneity here is allowed to affect policy *transmission*, but it plays no role in shaping policy *targets*, which remain concerned exclusively with macroeconomic aggregates.

For most of this section, we focus on the optimal setting of nominal interest rates i_t. We will only briefly refer to optimal stimulus check policy at the end.[18]

5.1 Problem set-up

We consider a policymaker with the traditional "dual mandate" objective function

$$\mathcal{L}^{DM} \equiv \sum_{t=0}^{\infty} \beta^t \left[\lambda_\pi \widehat{\pi}_t^2 + \lambda_y \widehat{y}_t^2\right]. \tag{22}$$

The policymaker wishes to stabilize inflation and output around the deterministic steady state, with weights λ_π and λ_y. A loss function of this form can be derived as a second-order approximation to household welfare in standard RANK models (Woodford, 2003), though we treat it as *ad hoc*; as noted, such a loss is a useful starting point because it is (i) of practical relevance, and (ii) expositionally useful, allowing us to disentangle the effects of heterogeneity on optimal policy design by transmission vs. targets. In sequence-space notation, this loss function becomes

$$\mathcal{L}^{DM} = \lambda_\pi \widehat{\boldsymbol{\pi}}' W \widehat{\boldsymbol{\pi}} + \lambda_y \widehat{\boldsymbol{y}}' W \widehat{\boldsymbol{y}}.$$

[18] To be more precise, we study the optimal monetary policy problem of setting i_t, with $\tau_{x,t} = 0$ in the background. There are, of course, still effects of fiscal adjustments through the endogenous fiscal rule (12).

The policymaker now sets nominal interest rates to minimize this dual-mandate loss (22) subject to the equilibrium constraints embedded in Definition 1. By Lemma 1, we can reduce these two constraints to two simple relationships: the Phillips curve (13) and the IS curve (14). This optimal policy problem is thus a minimal departure from optimal policy analysis in conventional representative-agent environments: the loss function (by assumption) and the supply side are entirely unaffected, while the demand constraint changes from a simple Euler equation as in (15) to the richer demand relation (14).

5.2 Optimal policy

Our analysis of optimal policy proceeds in two steps: we first present the optimal targeting rule, before then turning to practical implementation. We conclude with a very brief discussion on what happens if the policy instrument of choice were instead stimulus checks.

Optimal targeting rule. By the general arguments of Section 4, we obtain the optimal implicit targeting rule as

$$\lambda_\pi \Theta'_{\pi,i} W \widehat{\boldsymbol{\pi}} + \lambda_y \Theta'_{y,i} W \widehat{\boldsymbol{y}} = \mathbf{0}, \tag{23}$$

where $\Theta_{\pi,i}$ and $\Theta_{y,i}$ give the dynamic causal effects on $\boldsymbol{\pi}$ and $\boldsymbol{y}$, respectively, of a change in $\boldsymbol{i}$; for example, the first column in each case gives the impulse responses to a monetary policy change that only perturbs the nominal rate at date 0. It further follows from (23) that the optimal targeting rule in fact only depends on the *relative* effects of interest rate changes on inflation vs. output. The following proposition leverages this insight to re-write the optimal rule in a more instructive way.

Proposition 1. *Let $\widehat{\boldsymbol{c}}$ be a path of household consumption with zero net present value, i.e., $\sum_{t=0}^{\infty}\left(\frac{1}{1+\bar{r}}\right)^t \widehat{c}_t = 0$. If, for* any *such path $\widehat{\boldsymbol{c}}$, we have that*

$$\widehat{\boldsymbol{c}} \in \; image(\tilde{\mathcal{C}}_i), \tag{24}$$

then the optimal monetary policy rule for a dual-mandate policymaker with loss function (22) can be written as the forecast target criterion

$$\lambda_\pi \widehat{\pi}_t + \frac{\lambda_y}{\kappa}\left(\widehat{y}_t - \widehat{y}_{t-1}\right) = 0, \quad \forall t = 0, 1, \ldots \tag{25}$$

The key takeaway of Proposition 1 is that the optimal forecast target criterion, (25), is exactly the same in our economy as it is in standard representative-agent models, and thus invariant to household heterogeneity. The proposition begins with a mild regularity condition on $\tilde{\mathcal{C}}_i$, i.e. the mapping from nominal interest rate paths to net excess consumption demand in the generalized "IS" curve (14). That condition—readily verified analytically in simple "HANK-like" economies (Wolf, 2025) or numerically in richer economies like ours—simply states that nominal interest rate policy can flexibly manipulate intertemporal demand.[19] As soon as this is the case, the IS curve becomes a slack constraint in the policy problem; in particular, any desired path of output and inflation that is consistent with the Phillips curve can then be implemented through a suitable choice of interest rates. But since the Phillips curve as the supply side of the economy is (as per our assumptions) entirely independent of household heterogeneity, it follows that the optimal target criterion—and so also the optimal output-inflation allocation—is exactly as in RANK.

Implementation. While Proposition 1 establishes that optimal output and inflation outcomes will be unaffected by household heterogeneity, the switch from RANK to HANK may in principle still matter for the policy instrument path—i.e., for the time profile of nominal interest rates—necessary to implement this unchanged optimum.

In practice, however, even this implied policy instrument path is likely to be very close once we consider *quantitatively relevant* RANK and HANK models—i.e., model variants that are, like our model, consistent with evidence on monetary policy transmission. Intuitively, the optimal policy rule (23) yields an optimal path of output and inflation, and policy causal effects Θ_i allow us to recover the path of interest rates that implements this optimum. While household heterogeneity can in principle alter the Θ_i's and thus rate paths, the scope for this to change matters relative to RANK is limited once we compare HANK and RANK models that match empirical evidence on monetary shock transmission (i.e., parts of the Θ_i's). This observation is established in greater detail in Caravello et al. (2025).

Aside: optimal stimulus checks. Proposition 1 only considers the first instrument available to our policymaker: nominal interest rates. Results for stimulus checks follow immediately from Wolf (2025), who identifies conditions under which interest rate and stimulus check policies can implement the same sequences of aggregate output and inflation. More

[19]Formally, the condition says that, through manipulation of short-term nominal interest rates, the policymaker can engineer *any* possible net excess demand path with zero net present value. This is precisely what is needed for the IS curve to become a slack constraint.

formally, it follows from his results that, if

$$\widehat{\boldsymbol{c}} \in \text{image}(\tilde{\mathcal{C}}_x) \tag{26}$$

for all sequences $\widehat{\boldsymbol{c}}$ with zero net present value, then stimulus check policies can also implement the target criterion (25), just like conventional monetary policy. This alternative implementability condition (26) is again generally satisfied in HANK-type environments.

It follows from the previous discussion that the two policy instruments are perfect substitutes, and so that the solution to the problem of choosing both instruments to minimize (22) is indeterminate—multiple paths of the two policy instruments are consistent with the optimal outcomes for output and inflation. One way to break this indeterminacy is to introduce further constraints on instruments, e.g., a lower bound on nominal interest rates. The next section considers an alternative resolution to this indeterminacy: a richer loss function.

6 Adding distributional concerns

We now consider a policymaker that wishes to insulate individual households from the consequences of aggregate fluctuations. We begin by introducing the policymaker objective, and showing that the policy problem can be re-cast in the general linear-quadratic form of Section 4. We then characterize the optimal policy rule, and illustrate it through applications to distributional demand and cost-push shocks. We finally discuss how our policymaker compares to more standard welfare-based planners.

6.1 A distributional objective

This section introduces our policymaker objective. We begin by stating and providing an economic interpretation of the objective function. We then turn to our main result: optimal policy analysis based on this objective can be conveniently cast in linear-quadratic form.

THE POLICYMAKER OBJECTIVE. We consider a policymaker that wishes to insure households against business cycle-induced fluctuations in their (relative) consumption, but does *not* wish to provide additional insurance against idiosyncratic shocks, beyond what is already achieved through self insurance. We do so for two reasons. First, as our focus is on how *stabilization* policy should account for household heterogeneity, we want to abstract from any policymaker incentives to alter *long-run* allocations. In particular, we want to ensure that the policymaker does not seek to intervene in the absence of cyclical shocks, which is exactly what our assumptions will ensure. Second, it is an implication of our assumptions that the economy's deterministic steady state is efficient, thereby opening the door to linear-quadratic policy analysis (as in Woodford, 2003).[20] In the textbook representative-agent case, a production subsidy suffices for efficiency; in our environment, this is achieved through suitable time-varying planner weights, which precisely deliver a policymaker objective of the sort we consider, penalizing fluctuations in consumption shares.

To formally state our policymaker objective function, it will prove convenient to describe an individual's outcomes in terms of their history of idiosyncratic shocks; that is, we replace c_{it} with $\omega_t(\zeta_i^t)c_t$ where $\zeta_i^t \equiv (\zeta_{it}, \zeta_{it-1}, \zeta_{it-2}, \cdots)$ is individual i's history of idiosyncratic shocks and $\omega_t(\zeta_i^t) \equiv c_{it}/c_t$ is their share of aggregate consumption. Letting $\Gamma(\zeta)$ denote the (stationary) distribution of such histories (with ζ denoting a generic realization of a history), we will consider a social welfare function of the form

$$\mathcal{V} = \sum_{t=0}^{\infty} \beta^t \int \varphi(\zeta) \left[\frac{(\omega_t(\zeta)c_t)^{1-\gamma} - 1}{1-\gamma} - \nu(\ell_t) \right] d\Gamma(\zeta), \tag{27}$$

where $\varphi(\zeta)$ is a planner weight on the utility of households with history ζ. We note that this objective is an example of what Dávila & Schaab (2025) call a dynamic-stochastic planner.[21]

A SECOND-ORDER APPROXIMATION. Under suitable assumptions on the production subsidy and on the policymaker weights $\varphi(\zeta)$, we can ensure efficiency of our model economy's deterministic steady state; these assumptions are stated in Appendix D.3, and assumed from now on. Given efficiency, first-order terms in a second-order expansion of the policymaker loss vanish, and so we arrive at the following second-order representation:

[20] Our approach has some connection to the inverse optimal taxation literature (Heathcote & Tsujiyama, 2021; Le Grand et al., 2025). We require that the stationary equilibrium of the economy is optimal not just at date 0, but at all future dates as well. This in turn requires that the weight the planner attaches to an individual's welfare varies over time with their realization of idiosyncratic shocks, as we discuss further below. One possible interpretation of the time-varying weights is that they reflect unmodeled incentive concerns, limiting the extent of insurance the planner would like to provide.

[21] Dávila & Schaab (2025) decompose welfarist planner objectives into distinct considerations—aggregate efficiency, risk sharing, and redistribution. Dynamic-stochastic planner objectives arise when one only keeps some parts of the planner objective, but disregards others. Our objective function incorporates aggregate efficiency and a desire to provide insurance against business cycle-induced fluctuations in the consumption distribution, but does not value insurance against idiosyncratic shocks. As a result, there could be Pareto-improving trades (related to such idiosyncratic risk) that our policymaker simply does not value.

Proposition 2. *To second order, the social welfare function $\mathcal{V}$ is proportional to $-\mathcal{L}$,with*

$$\mathcal{L} \equiv \sum_{t=0}^{\infty} \beta^t \left[\widehat{\pi}_t^2 + \lambda_y \widehat{y}_t^2 + \lambda_\omega \int \frac{\widehat{\omega}_t(\zeta)^2}{\bar{\omega}(\zeta)} d\Gamma(\zeta) \right], \tag{28}$$

where $\widehat{\omega}_t(\zeta) \equiv \omega_t(\zeta) - \bar{\omega}(\zeta)$, $\bar{\omega}(\zeta)$ is the steady-state consumption share of an individual with history ζ, and λ_y and λ_ω are composite parameters defined in Appendix D.3.

Note that, in the representative-agent analogue of our economy, the loss function would feature the same first two terms, as is already well-known from prior work (e.g., see Woodford, 2003).[22] Household heterogeneity then simply adds a third, inequality-related term, with the planner wishing to stabilize the consumption *shares* of everyone in the economy.[23]

In the next few subsections we will analyze optimal policy outcomes under a loss function of the form (28). After doing so, we will in Section 6.5 discuss further how our planner loss relates to a welfare-based objective function.

6.2 Policy problem and optimal rule

We now cast the optimal policy problem of our distributional policymaker in the general linear-quadratic form of Section 4. Several further details, in particular on computation, are relegated to Appendix C.2.

Policy problem. Moving to a sequence-space formulation, we can write the policymaker loss as

$$\mathcal{L} = \widehat{\boldsymbol{\pi}}' W \widehat{\boldsymbol{\pi}} + \lambda_y \widehat{\boldsymbol{y}}' W \widehat{\boldsymbol{y}} + \int \lambda_\zeta \widehat{\boldsymbol{\omega}}(\zeta)' W \widehat{\boldsymbol{\omega}}(\zeta) d\Gamma(\zeta), \tag{29}$$

where $\lambda_\zeta \equiv \lambda_\omega / \bar{\omega}(\zeta)$. The consumption share for each idiosyncratic household history thus emerges as a separate target variable for the policymaker, in addition to the usual concern

[22] The fact that the first two terms in (28) are exactly the same as in the textbook RANK model depends on the specific assumptions that we make about the production side of our model, notably the aggregation of the supply block to a simple NKPC. In more general HANK environments these first two terms—i.e., the representative-agent component of the loss—would differ. The shape of the inequality term, in contrast, is more general, depending only on our assumptions on consumer preferences and equal labor rationing.

[23] Bilbiie (2008) and Acharya et al. (2023) derive quadratic loss functions in models with limited heterogeneity. In those cases, the inequality term is a function of aggregate output, so the concern for inequality manifests as a larger weight on output stabilization. The optimal policy rule then takes the form of a simple target criterion similar to (25), but with appropriately adjusted coefficients.

for aggregate inflation and output.[24]

The constraints of the optimal policy problem characterize the evolution of policymaker targets—π, y, and the consumption shares $\omega(\zeta)$—as a function of exogenous shocks and policy choices. As we discussed in Section 5, the evolution of aggregate output and inflation are governed by the Phillips curve (13) and the IS curve (14). It then remains to describe the evolution of the inequality term in (28). In Appendix C.2, we establish that, to first order, we can write the consumption share for a household with specific history ζ as

$$\widehat{\boldsymbol{\omega}}(\zeta) = \Omega_{\omega(\zeta),x}\widehat{\boldsymbol{x}}, \qquad \forall\zeta \tag{30}$$

where $\boldsymbol{x} = (\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\tau}_x, \boldsymbol{\tau_e}, \boldsymbol{m}, \pi_0, q_0^k, q_0^b)$ stacks the prices and shocks that affect the household decision problem, and the maps $\Omega_{\omega(\zeta),x}$ give the derivatives of consumption shares with respect to these aggregate variables. The intuition for (30) is the same as that for the aggregate consumption function—an individual household's consumption, given their history of idiosyncratic shocks, evolves over time as a function of the aggregate inputs to the household consumption-savings problem. Finally, echoing the discussion in Section 4, we again re-write this constraint in impulse response space, solving out the dependence of consumption shares on endogenous aggregates:

$$\widehat{\boldsymbol{\omega}}(\zeta) = \Theta_{\omega(\zeta),i}\widehat{\boldsymbol{i}} + \Theta_{\omega(\zeta),\tau_x}\widehat{\boldsymbol{\tau}}_x + \Theta_{\omega(\zeta),\eta}\widehat{\boldsymbol{\eta}} + \Theta_{\omega(\zeta),m}\widehat{\boldsymbol{m}}. \qquad \forall\zeta \tag{31}$$

This concludes the set-up of the linear-quadratic problem.

Optimal policy. Having expressed the optimal policy problem in linear-quadratic form, we can finally apply the general results of Section 4 to provide a characterization of optimal policy rules. For monetary policy we find

$$\Theta'_{\pi,i}W\widehat{\boldsymbol{\pi}} + \lambda_y\Theta'_{y,i}W\widehat{\boldsymbol{y}} + \int \lambda_\zeta\Theta'_{\omega(\zeta),i}W\widehat{\boldsymbol{\omega}}(\zeta)d\Gamma(\zeta) = \mathbf{0}, \tag{32}$$

where as before the matrices $\Theta_{\bullet,i}$ collect the causal effects of nominal interest rate movements on the various policymaker targets. In particular, we can see that the first two terms in (32) are identical to the optimal dual-mandate rule (23), just now with weights derived from

[24] Note that, technically, (29) does not immediately fit into our framework in Section 4 since the objective here features an integral (rather than a simple sum). Our approach to computation uses an equivalent formulation of the problem with finitely many policy targets (see Appendix C.2).

policymaker preferences, rather than exogenously assumed. The novel third term—which reflects the planner's distributional insurance concerns—collects the causal effects of interest rate movements on consumption shares. Proceeding analogously for stimulus checks we find the optimal fiscal rule

$$\Theta'_{\pi,\tau_x} W\widehat{\boldsymbol{\pi}} + \lambda_y \Theta'_{y,\tau_x} W\widehat{\boldsymbol{y}} + \int \lambda_\zeta \Theta'_{\omega(\zeta),\tau_x} W\widehat{\boldsymbol{\omega}}(\zeta) d\Gamma(\zeta) = \mathbf{0}. \tag{33}$$

As discussed in Section 5, nominal interest rate and stimulus check policy in our environment can implement exactly the same output-inflation allocations, so the first two terms in (32) and (33) will reflect identical aggregate stabilization objectives. The distributional terms $\Theta_{\omega(\zeta),i}$ and $\Theta_{\omega(\zeta),\tau_x}$, on the other hand, will generally differ, thus opening the door for interest rate and stimulus check policies to be useful complementary tools for aggregate stabilization. Overall, (32) and (33) fully characterize joint optimal monetary-fiscal policy.

A key takeaway from those two forecast targeting rules is that optimal policy deviates from the dual-mandate case—i.e., is shaped by household heterogeneity—if and only if the policy instruments can affect cross-sectional consumption inequality. On the one hand, if monetary policy is distributionally neutral (as is the case in the environment of Werning, 2015, see Appendix A.2 for details) then household inequality does not affect the optimal rule: setting $\Theta_{\omega(\zeta),i} = \mathbf{0}$ and simplifying (32), we find that the optimal rule continues to take a standard dual-mandate form[25]

$$\widehat{\pi}_t + \frac{\lambda_y}{\kappa} (\widehat{y}_t - \widehat{y}_{t-1}) = 0, \quad \forall t = 0, 1, \dots \tag{34}$$

In this case, monetary policy optimally focuses on aggregate objectives while fiscal policy instead takes sole responsibility for all distributional considerations. On the other hand, if interest rate cuts are strongly progressive (as is the case in Bhandari et al., 2021; Dávila & Schaab, 2022), then the concerns about inequality as embedded in the policymaker objective (28) may materially change optimal monetary policy conduct, with insurance concerns possibly swamping the usual price and output stabilization motives.

To summarize, the analysis in this section has revealed that the extent to which household inequality shapes optimal policy depends, first, on whether there are distributional

[25] Bilbiie (2024) makes an analogous argument in a two-agent context, showing that, if monetary policy does not redistribute between spenders and savers, then the optimal policy rule is not affected by inequality concerns. Our analysis reveals that the conditions underlying Werning's aggregation result are precisely enough to extend this insight to our heterogeneous-agent setting.

considerations the policymaker would like to address, and second, on whether, through the use of her instruments, she is in a position to address them. Our expanded policymaker objective opens the door for the first margin. The model calibration results and underlying empirical evidence in Section 3, however, suggest that, at least for monetary policy, the second necessary ingredient is largely absent in our model—monetary policy is simply an unattractive tool for achieving distributional objectives. Our quantitative experiments in the next two subsections substantiate this intuition through applications to particular macroeconomic shocks.

6.3 A distributional shock

We use our calibrated model as a laboratory to explore optimal policy following an aggregate shock with strong distributional consequences, depressing consumption of the poor relative to that of the rich, somewhat akin to the Covid-19 recession. Formally, the shock is an innovation to m_t, the exogenous driver of income dispersion in the income process (16). This shock redistributes labor income from low-income to high-income households, depressing aggregate demand—spending falls as income is redistributed towards lower-MPC households. We then study three optimal policy responses to this aggregate demand shock: dual-mandate monetary policy, monetary policy with distributional targets, and finally joint monetary and transfer policy with distributional targets. All results are displayed in Figure 2.

Dual mandate. We begin with the optimal dual-mandate policy response (in grey). As aggregate demand falls following the distributional shock, the dual-mandate central banker cuts nominal interest rates to perfectly stabilize output and inflation (i.e., she attains the classic "divine coincidence"). The bottom panel shows how the consumption distribution changes on impact of the shock: the shock itself redistributes from low-income to high-income households, and monetary policy does little to offset that.

Monetary policy with distributional concerns. We next consider the optimal monetary policy for a policymaker that also has distributional objectives (in orange-dashed). Our headline finding here is that the optimal policy response is very similar to the dual-mandate policy, with nominal interest rates now cut somewhat more on impact of the shock but then very closely following the dual mandate policy. As a result, output and inflation continue to be stabilized almost perfectly. Relative to the dual-mandate case, all households consume slightly more in the initial period, but the shape of the consumption distribution is

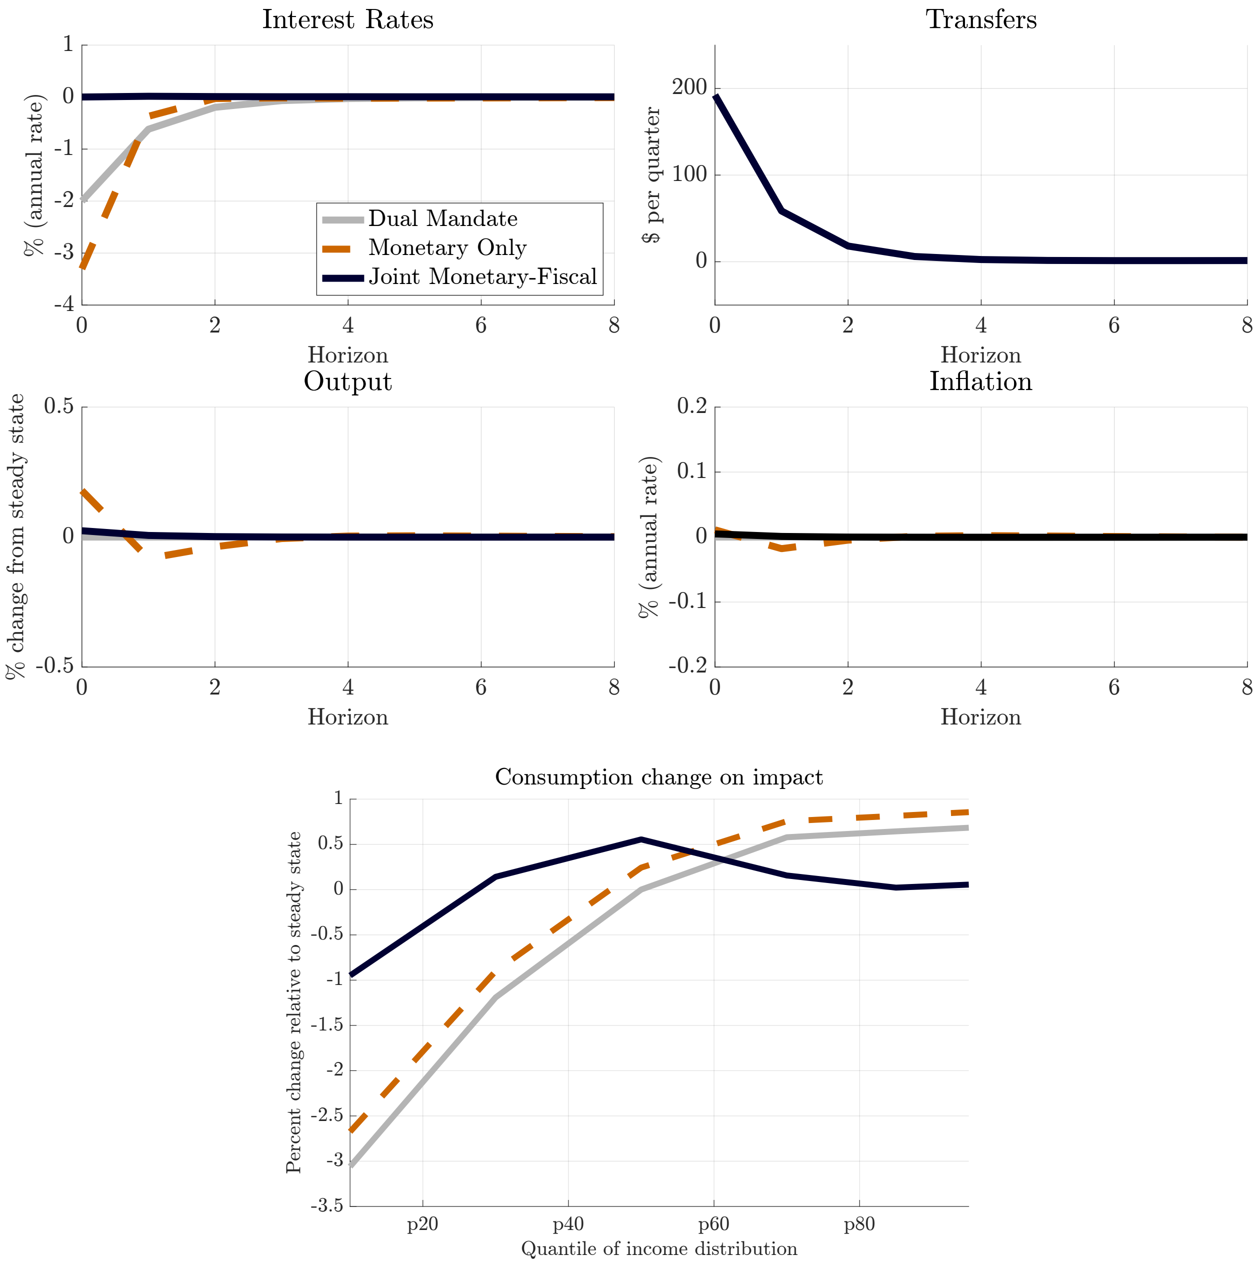

Figure 2: Optimal policy response to an income distribution shock m_t. The figure shows the results from three policy rules: optimal monetary policy for a dual-mandate policymaker (grey), optimal monetary policy with distributional targets (orange dashed), and optimal policy for a planner that can use both tools (dark blue). Transfers are expressed in units of dollars, using the conversion that steady state per capita GDP is $60,000.

hardly altered from the dual-mandate outcome. The intuition for this is exactly as already anticipated at the end of Section 6.2. The original shock m_t has a strong distributional tilt, with low-income households losing the most. The policymaker would like to lean against this redistribution and stabilize consumption shares. However, as shown in Figure 1, monetary policy in our model can only raise the consumption of all households more or less uniformly. It thus in particular follows that a policy that stabilizes consumption at the bottom at the same time over-stimulates consumption at the top, leading to substantial aggregate overheating. We provide a quantitative illustration of this point in Figure B.6, where we simulate a policy rule that targets the consumption of households with low income and wealth.

Joint policy. Finally we turn to the optimal policy that uses both nominal interest rates and stimulus checks to jointly pursue aggregate and distributional objectives (in dark blue). Compared to monetary policy, fiscal stimulus payments are much more progressive, with the consumption of low-income and low-wealth households responding significantly more than that of high-income, high-wealth households (see Figure B.5). Intuitively, this reflects both differences in MPCs as well as differences in income levels, with any given dollar amount of (uniform) stimulus checks amounting to a much larger fraction of income at the bottom of the income distribution. As a result, fiscal stimulus payments are particularly well-suited as a tool to address shocks that differentially affect low-income and high-income households. The results depicted in Figure 2 are consistent with this intuition: in response to the distributional shock m_t, the policymaker can use the checks to almost perfectly stabilize aggregate output, inflation, *and* consumption inequality. Since stimulus payments both compress inequality and stimulate aggregate demand, there is little need to adjust interest rates.

Discussion. Overall, our analysis of optimal policy responses to the distributional demand shock m_t suggests two broad lessons for optimal cyclical policy design. First, recall from our prior analysis that a monetary policymaker with distributional targets has an incentive to deviate from the prescriptions of dual mandate policy if and only if monetary policy can help her to offset the distributional incidence of the underlying shock. With our calibration yielding a relatively uniform distributional incidence of monetary interventions, it follows that monetary policy is ill-suited as a tool to attain distributional objectives. In particular, stabilizing consumption at the bottom of the distribution may require significant departures from dual-mandate objectives and so aggregate efficiency; if the policymaker is unwilling to accept these departures (as was the case here), then the optimal policy will tend to stay rather close to the dual-mandate benchmark, as seen here. Second, since fiscal stimulus payments and nominal interest rate policy have very different cross-sectional incidence profiles, they can actually be highly complementary policy tools. For the shock considered here, fiscal transfers were a particularly well-suited tool; for shocks with uniform incidence, the policymaker may instead want to mostly rely on monetary policy; and for other shocks, a mixture of the two instruments may be desirable.

6.4 A cost-push shock

Our second shock is an inflationary cost-push shock, showing up as a wedge η_t in the aggregate Phillips curve. As usual, this shock induces an output-inflation trade-off. In response, a dual-mandate central bank persistently increases interest rates (after a short-lived drop at the beginning), depressing output and moderating inflationary pressures, all displayed in grey in the top panels of Figure 3. Since the incidence of monetary policy is quite close to uniform, and since (by our assumptions) the labor share is fixed, it follows that consumption declines roughly uniformly on impact, in tandem with the decline in output. A central banker with distributional objectives thus sees no need to deviate from the dual-mandate outcomes; accordingly, in Figure 3, the orange-dashed lines (for the policymaker with distributional targets) are close to the dual-mandate outcomes (grey).[26]

To gain further insight into our results and their relation to prior work, we now compare our findings with those of an analogous experiment in the sticky-*price* version of our model, as introduced in Section 3.2 (and discussed further in Appendix B.3). Specifically, in Figure 4, we show what happens to aggregate output and the consumption distribution following a cost-push shock in that model. The results are shaped by two key properties of the sticky-price model. First, monetary policy is now progressive rather than distributionally neutral: interest rate cuts now boost real wages and depress mark-ups, thus disproportionately increasing consumption of those that rely on labor income rather than capital income (see Figure B.4). Second, a contractionary cost-push shock—which increases mark-ups—redistributes from low-wealth (and low-income) households to the upper end of the distribution. Putting the pieces together, we see that the dual-mandate policy response contracts output (as usual) and chiefly hurts the poor (grey). A monetary policymaker with distributional concerns not only dislikes that distributional tilt, but also can do something about

[26]The cost-push shock is thus different in nature from the distributional shock studied in Section 6.3. There, the distributional policymaker would have liked to redistribute, but finds monetary policy limited as a tool for that purpose. Here, instead, she does not even *want* to redistribute, simply because the shock itself does not have a strong distributional tilt.

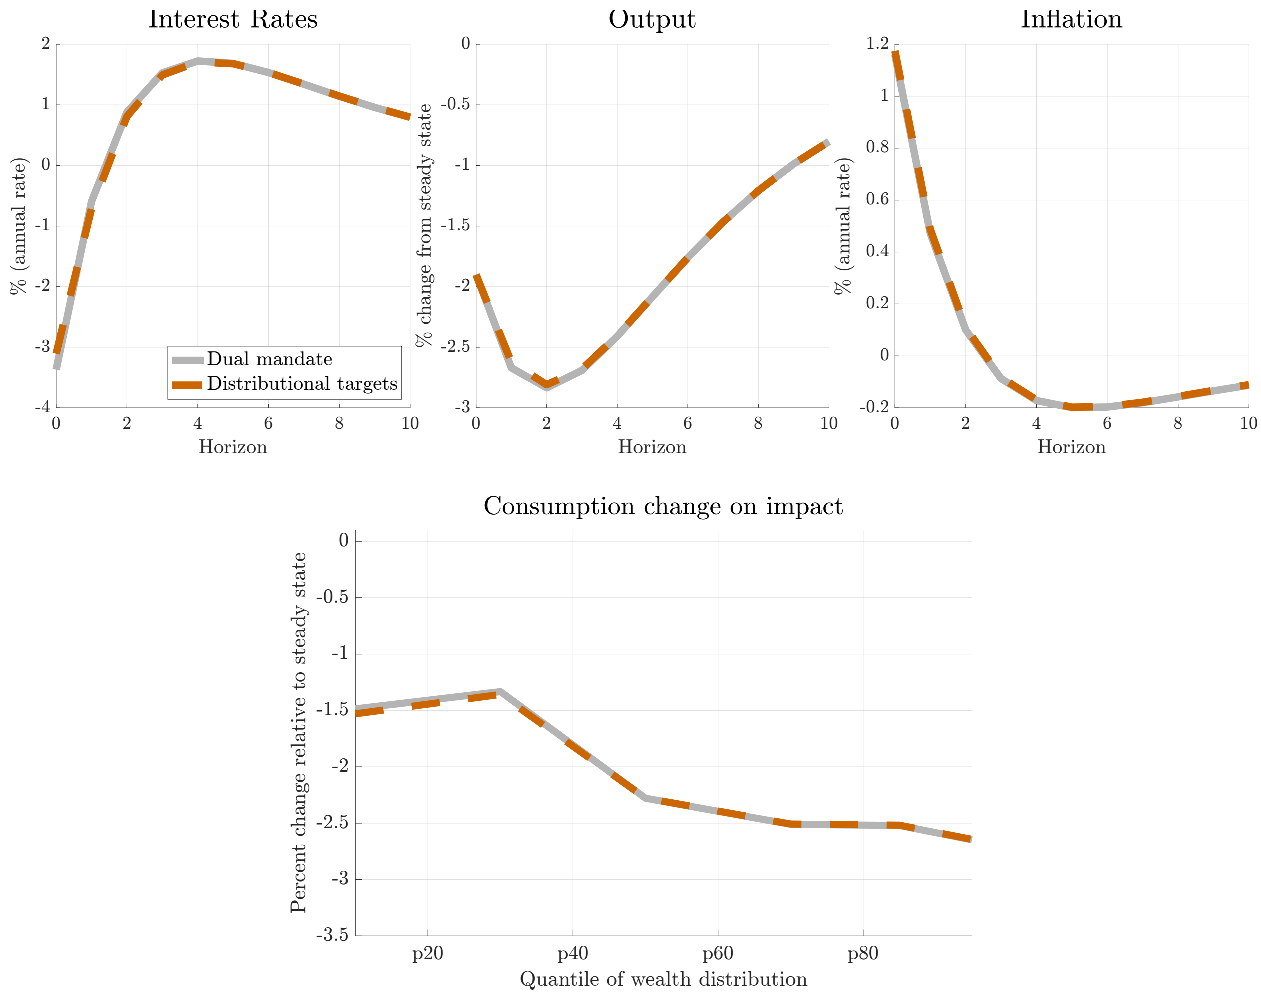

Figure 3: Optimal policy response to an inflationary cost-push shock. The figure shows optimal monetary policy for a policymaker whose objective is to stabilize output and inflation (grey) and for a policymaker with distributional targets (orange dashed).

it, by depressing output less aggressively (orange-dashed). While this model implies that distributional concerns have an important impact on policy, it relies on a counterfactually strong effect of policy on the labor share (see Figure B.3).

Contrasting Figures 2 and 3 with Figure 4, we see how distributional concerns in principle *can* shape optimal policy design—it is just a feature of our model environment that they do not. In our model, monetary policy only has a weak distributional tilt, and so outcomes are close to the dual-mandate benchmark, irrespective of whether the policymaker would *like* to change the distributional outcomes (as in Figure 2) or not (as in Figure 3). Switching to a model that, along key dimensions, makes counterfactual predictions about the distributional effects of monetary policy, overturns those results, and optimal distributional policy starts

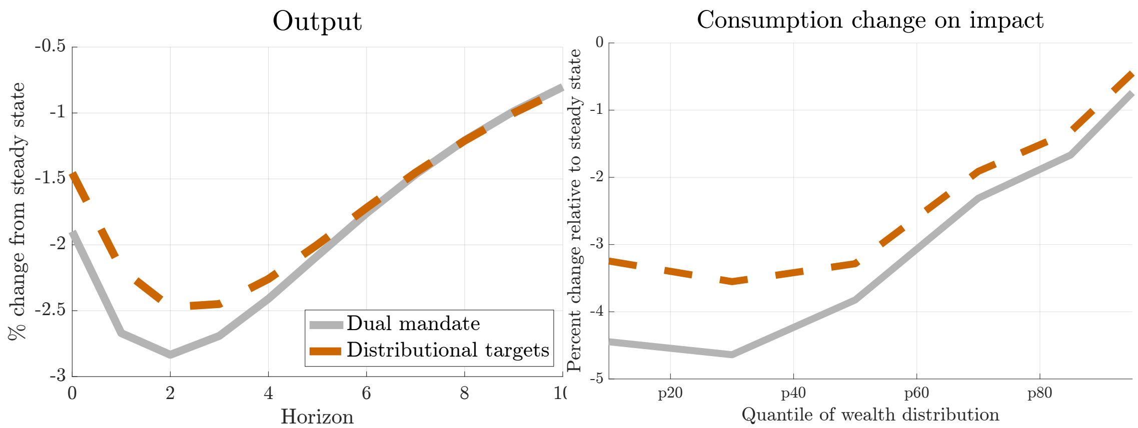

Figure 4: Optimal policy response to an inflationary cost-push shock in sticky-price model. The figure shows optimal monetary policy for a policymaker whose objective is to stabilize output and inflation (grey) and for a policymaker with distributional targets (orange dashed).

to deviate more from the dual-mandate benchmark.

6.5 Comparison with a welfare-based planner

The fundamental difference between our policymaker objective function and conventional welfare-based planner objectives is that, in the latter, the planner weights that are assigned to individual households are fixed, rather than dynamic and stochastic. Recall that we can express our baseline planner objective as

$$\mathcal{V} = \int \left\{ \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \varphi_{it} \left[\frac{c_{it}^{1-\gamma}}{1-\gamma} - v(\ell_{it}) \right] \right\} di, \tag{35}$$

where φ_{it} is a function of the history of the idiosyncratic shocks experienced by household i up to and including date t. We so far ensured, through time variation in the φ_{it}'s, that our economy's steady state was efficient, so the planner would only want to offset the distributional effects of *shocks*. As a point of comparison, we will now consider a more conventional alternative, which we denote $\mathcal{V}^W$, that takes the same form as (35), but now with φ_{it} replaced by a constant value $\bar{\varphi}_i$. Specifically, we set $\bar{\varphi}_i = \varphi_{i0}$ so the $\mathcal{V}^W$ planner has no desire to change the date-0 consumption allocation; afterwards, however, since the weights remain fixed (rather than varying stochastically), she would in general desire redistribution even in the absence of any aggregate shocks.

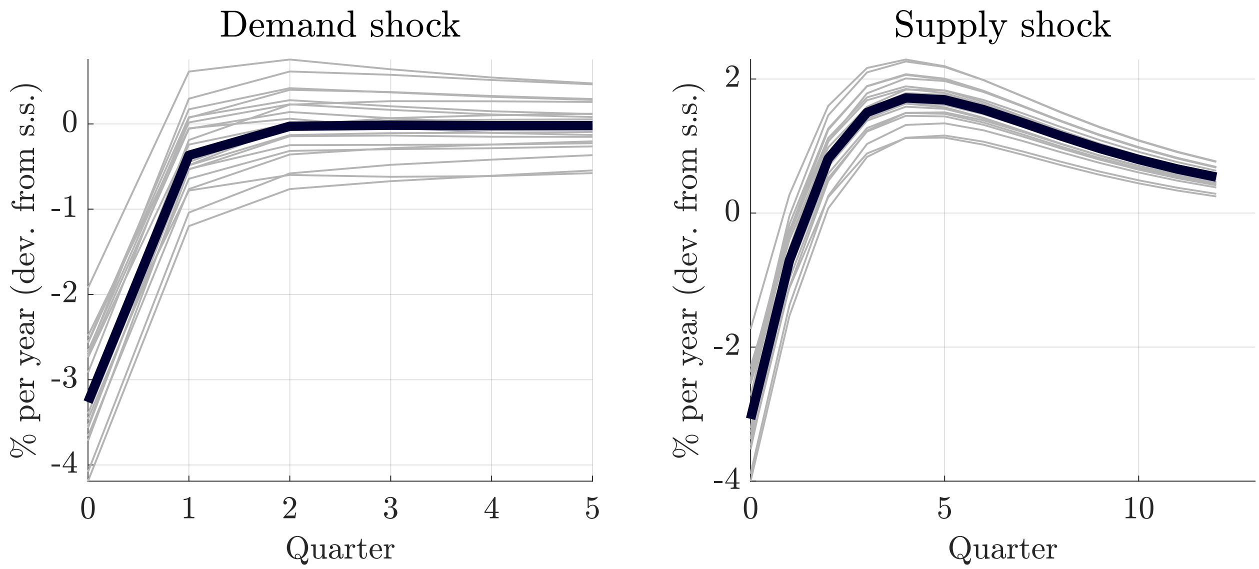

Figure 5: Simulated policy paths (nominal interest rates) around the optimal policy for our baseline policymaker objective $\mathcal{V}$ (bold line).

Comparing planners—the exercise. In the preceding sections we studied optimal policy following two kinds of business-cycle shocks: the demand shock m_t and the supply shock η_t. The bold lines in Figure 5 repeat the results of these exercises, showing the optimal monetary policy responses (i.e., nominal interest rate paths) to those shocks for our baseline policymaker with objective function $\mathcal{V}$. We now wish to evaluate whether the alternative planner with the objective $\mathcal{V}^W$ agrees with our baseline planner in the ranking of these "optimal" policy options relative to several other, nearby alternatives.

Specifically, we construct a randomly selected set of paths for the nominal interest rate in the vicinity of the optimum (shown as the grey lines in Figure 5), and then for each such path compute allocations—i.e., paths of macroeconomic aggregates and from there household outcomes, giving us a panel of $\{c_{it}, \ell_t\}_{t=0}^T$ for all households i. We then compute $\mathcal{V}^W$ and $\mathcal{V}$ for each path and assess how they compare.[27] By design, the optimal path gives the highest value for the objective $\mathcal{V}$, but we want to evaluate whether the same is true for $\mathcal{V}^W$. Consistently with our focus on business-cycle fluctuations we evaluate welfare over the first

[27] The interest rate paths are constructed as $i_t = i_t^* + \nu_1 0.9^t + \nu_2 0.6^t + \nu_3 0.3^t$, where i_t^* is the optimal policy and each ν_j is drawn from $N(0.0004, 0.001^2)$. Given the sampled interest rate path, we then solve for aggregate prices, and finally simulate a panel of households. Across each simulation, we hold fixed the sequence of idiosyncratic shocks that we use for each household; this implies that household i at date t is associated with a particular history of idiosyncratic shocks across each simulation.

ten years following the initial shock.[28]

Results. The top panels of Figure 6 show the main findings of our analysis. The panels show policymaker objectives (y-axis) for different interest rate paths (x-axis), for the demand shock (left panel) and the supply shock (right panel), in thick blue for the baseline objective, and in thin orange for the alternative objective. The simulated policy paths (on the x-axis) are ordered from worst to best according to the baseline policy objective $\mathcal{V}$, so the thick blue lines are increasing by construction, with the optimal policy on the right. The key finding is that the thin orange lines are also generally upward-sloping, i.e., that the policy paths preferred under $\mathcal{V}$ are generally also preferred under $\mathcal{V}^W$. We see that this is true for both demand and supply shocks.

To understand where those results are coming from, we in the middle and bottom panels decompose the policy objective into two components: aggregate efficiency and distributional concerns. Normalizing planner weights so that $\int \varphi_{it} di = 1 \forall t$, we can rearrange (35) as[29]

$$\mathcal{V} = \underbrace{\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{\Xi c_t^{1-\gamma}}{1-\gamma} - v(\ell_t) \right]}_{\text{aggregate efficiency}} + \underbrace{\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\int \varphi_{it} \left(\frac{c_{it}^{1-\gamma}}{1-\gamma} - \frac{\Xi c_t^{1-\gamma}}{1-\gamma} \right) di \right]}_{\text{distributional}}. \tag{36}$$

The same kind of decomposition applies for $\mathcal{V}^W$, with φ_{it} replaced by $\bar{\varphi}_i$. By construction, the two objective functions agree on aggregate efficiency, as displayed in the middle panel. On the distributional term, however, the two objective functions do somewhat differ (bottom panel), simply because the weights here are different.[30] Our main result—that the *overall* welfare assessments are nevertheless similar—is for two main reasons. First, and most importantly, the distributional effects of monetary policy are relatively small in our model. It follows that, even if the two objectives were to disagree in their evaluation of how consumption is

[28] Over longer frequencies, the persistent movements in the consumption distribution caused by even transitory changes in monetary policy further increase differences in the distributional term in Figure 6 (similar to the slow-moving dynamics in Bayer et al., 2024). Accumulating these effects over sufficiently long horizons can lead to some (still generally modest) differences in the rankings of the two objective functions.

[29] We use $\ell_{it} = \ell_t \forall i$. The constant Ξ simply scales the consumption utility of the fictitious representative household to account for dispersion in the marginal utilities of households; see (D.16).

[30] The disagreement between $\mathcal{V}^W$ and $\mathcal{V}$ reflects their different motivations—$\mathcal{V}$ wishes to push the consumption distribution back to what it would be without aggregate shocks, while $\mathcal{V}^W$ wishes to push it back to what it would have been *in date 0* without shocks. Even in the absence of any aggregate shocks, $\mathcal{V}^W$ views it as undesirable that households with high consumption at date 0 gradually spend down their wealth and end up consuming less. Different paths for the macroeconomy can offset or amplify such reversals, which $\mathcal{V}^W$ likes or dislikes, respectively.

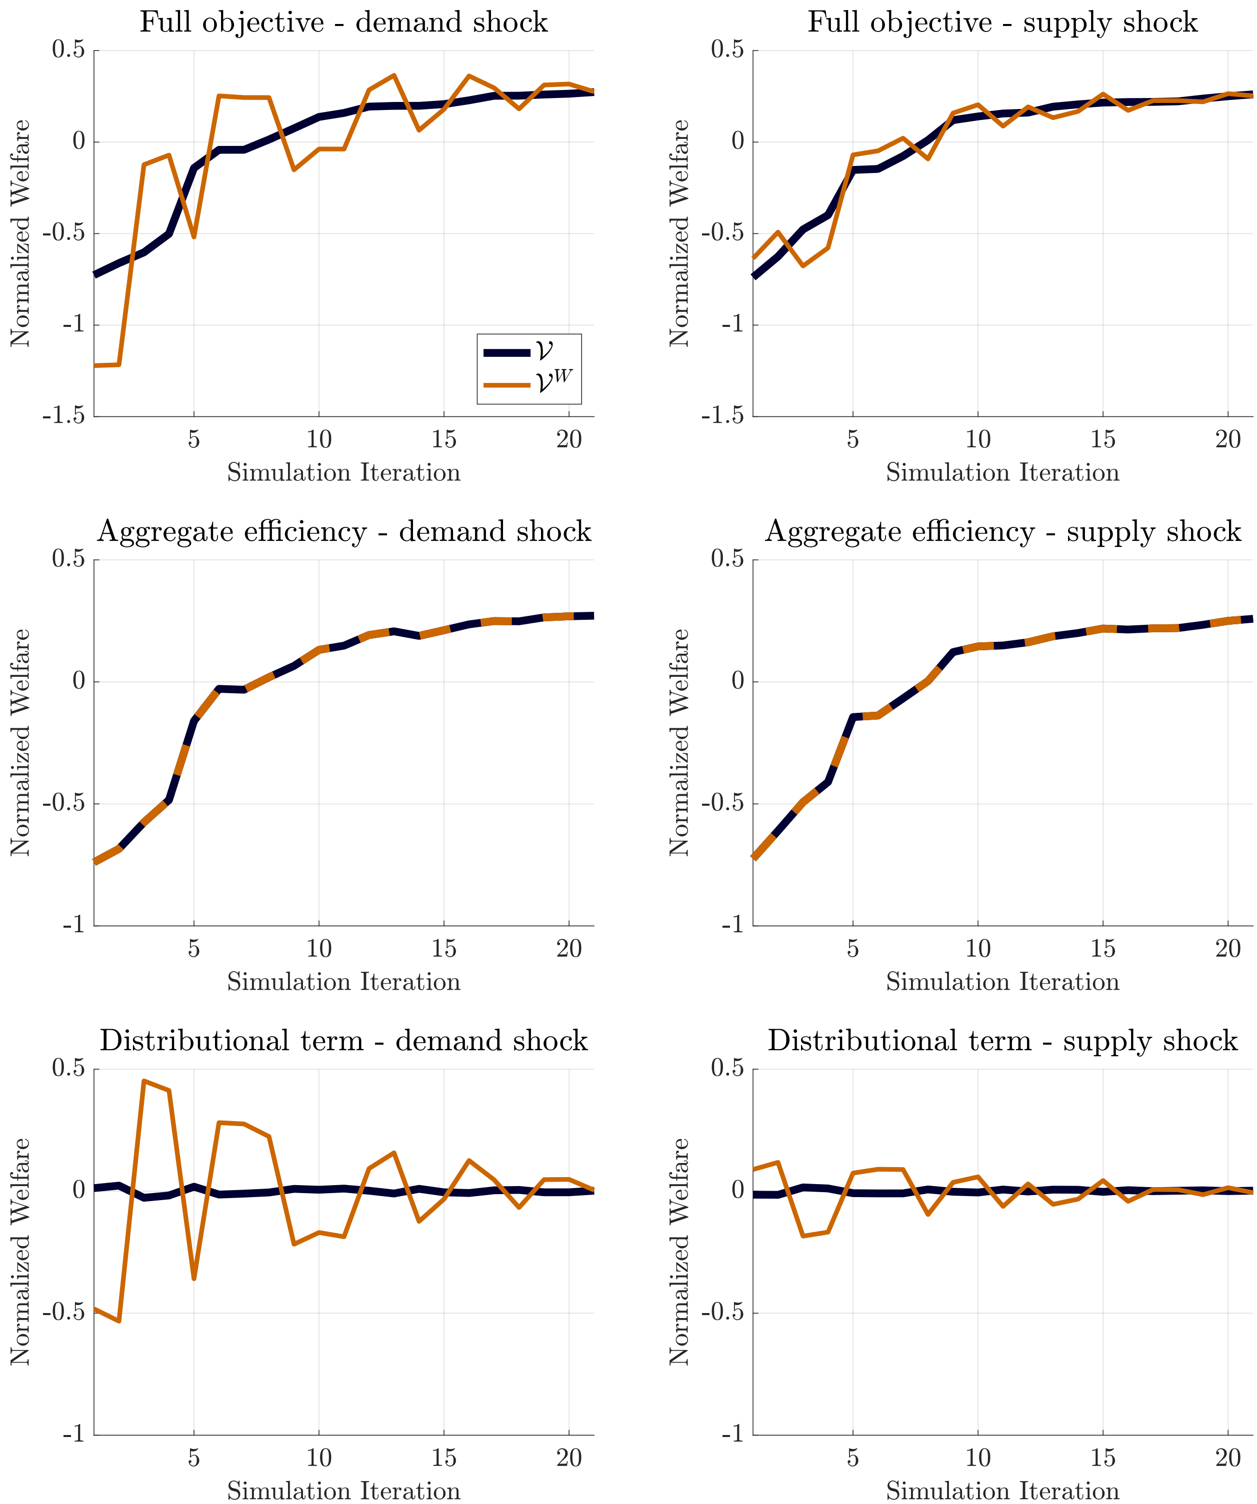

Figure 6: Rankings of policy paths according to $\mathcal{V}$ and $\mathcal{V}^W$. The simulated policy paths are ordered according to $\mathcal{V}$ with worst on the left and best on the right. For each objective function and shock we scale the results so that $\mathcal{V}_{\text{best}} - \mathcal{V}_{\text{worst}} = 1$. The lower panels give an additive decomposition of the top panel according to (36). $\mathcal{V}$ and $\mathcal{V}^W$ are calculated using a simulated panel of households of 100,000 individuals over 40 quarters following the shock.

distributed, changing monetary policy does not much affect that distribution, and so—under both objectives—the ranking of different monetary policy options is largely dominated by aggregate efficiency concerns. Second, the macroeconomic shocks we study are relatively transitory while idiosyncratic productivity is relatively persistent, and so the weights in our baseline objective function $\mathcal{V}$ do not change too much over business-cycle frequencies.

7 Conclusion

Should household inequality affect the conduct of cyclical stabilization policy? The analysis in this paper suggests the following three main takeaways.

First, for central banks that target standard economic aggregates (i.e., a traditional "dual mandate"), household inequality—at least as typically modeled in the HANK literature—is likely to only have modest effects. Analytically, we have given conditions under which the optimal forecast target criterion of a dual-mandate central banker is unaffected by household inequality, implying that aggregate output and inflation outcomes in response to any shock will be exactly as in analogous representative-agent environments. Empirically, a long literature already estimates the causal effects of interest rate changes on output and inflation. Models that are consistent with this evidence will necessarily yield similar predictions for the paths of rates necessary to implement a given output and inflation target—irrespective of whether the model features household heterogeneity or not. We stress, however, that our analysis was premised on a view that such heterogeneity only affects demand, and leaves the supply side of the economy unaffected; one avenue for future work is to relax that assumption.

Second, the extent to which distributional objectives shape optimal monetary policy design depends crucially on the causal effects of nominal interest rate changes on consumption inequality. According to our analysis, monetary policy has relatively even effects on consumption across households. Interest rate policy is thus not a particularly sharp tool to deal with shocks that disproportionately affect the poor, at least not without substantial costs in terms of aggregate stabilization.

Third, stimulus checks—an alternative tool of stabilization policy, used increasingly frequently in recent decades—may be a more useful tool for distributional purposes. Such stimulus checks achieve stabilization *through* insurance at the bottom, thus making them well-suited to address cyclical fluctuations that disproportionately affect poor households. Of course, in practice, these advantages of fiscal stimulus checks must be weighed against the classic challenges in using fiscal policy as a stabilization tool described by Friedman (1968).

Data Availability Statement

The data codes underlying this article are available in Zenodo, at https://dx.doi.org/10.5281/zenodo.202148 The datasets were derived from sources in the public domain. Please see the replication package for a full list of datasets and citations.

References

Acharya, S., Challe, E., & Dogra, K. (2023). Optimal monetary policy according to HANK. *American Economic Review*, *113*(7), 1741–82.

Alves, F., Kaplan, G., Moll, B., & Violante, G. L. (2020). A further look at the propagation of monetary policy shocks in HANK. *Journal of Money, Credit and Banking*, *52*(S2), 521–559.

Amberg, N., Jansson, T., Klein, M., & Picco, A. R. (2022). Five facts about the distributional income effects of monetary policy shocks. *American Economic Review: Insights*, *4*(3), 289–304.

Andersen, A. L., Johannesen, N., Jørgensen, M., & Peydró, J.-L. (2023). Monetary policy and inequality. *The Journal of Finance*, *78*(5), 2945–2989.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, *109*(6), 2333–67.

Auclert, A., Bardóczy, B., & Rognlie, M. (2023). MPCs, MPEs, and multipliers: A trilemma for new keynesian models. *The Review of Economics and Statistics*, *105*(3), 700–712.

Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, *89*(5), 2375–2408.

Auclert, A., Rognlie, M., Souchier, M., & Straub, L. (2021). Exchange rates and monetary policy with heterogeneous agents: Sizing up the real income channel. Technical report, National Bureau of Economic Research.

Auclert, A., Rognlie, M., & Straub, L. (2024). The intertemporal keynesian cross. *Journal of Political Economy*, *132*(12), 4068–4121.

Auclert, A., Rognlie, M., & Straub, L. (2025). Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics*, *17*.

Basu, S. & Fernald, J. G. (1997). Returns to scale in us production: Estimates and implications. *Journal of Political Economy*, *105*(2), 249–283.

Bayer, C., Born, B., & Luetticke, R. (2024). Shocks, frictions, and inequality in us business cycles. *American Economic Review*, *114*(5), 1211–1247.

Benigno, P. & Woodford, M. (2012). Linear-quadratic approximation of optimal policy problems. *Journal of Economic Theory*, *147*(1), 1–42.

Bhandari, A., Evans, D., Golosov, M., & Sargent, T. J. (2021). Inequality, business cycles, and monetary-fiscal policy. *Econometrica*, *89*(6), 2559–2599.

Bilbiie, F. O. (2008). Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of economic theory*, *140*(1), 162–196.

Bilbiie, F. O. (2024). Monetary policy and heterogeneity: An analytical framework. *Review of Economic Studies*, *92*(4), 2398–2436.

Bilbiie, F. O., Monacelli, T., & Perotti, R. (2024). Stabilization vs. redistribution: The optimal monetary–fiscal mix. *Journal of Monetary Economics*, *147*.

Boar, C. & Midrigan, V. (2022). Efficient redistribution. *Journal of Monetary Economics*, *131*, 78–91.

Boppart, T., Krusell, P., & Mitman, K. (2018). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, *89*, 68–92.

Broer, T., Kramer, J. V., & Mitman, K. (2022). The curious incidence of monetary policy across the income distribution. Technical report, Sveriges Riksbank Working Paper Series.

Caravello, T. E., McKay, A., & Wolf, C. (2025). Evaluating monetary policy counterfactuals: (when) do we need structural models? MIT working paper.

Castaneda, A., Diaz-Gimenez, J., & Rios-Rull, J.-V. (2003). Accounting for the us earnings and wealth inequality. *Journal of Political Economy*, *111*(4), 818–857.

Chang, M. & Schorfheide, F. (2022). On the effects of monetary policy shocks on earnings and consumption heterogeneity. CEPR Discussion Paper No. DP17049.

Coibion, O., Gorodnichenko, Y., Kueng, L., & Silvia, J. (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics*, *88*, 70–89.

Congressional Budget Office (2019). Projected changes in the distribution of household income, 2016 to 2021. Technical report, Congressional Budget Office. Report.

Dávila, E. & Schaab, A. (2022). Optimal monetary policy with heterogeneous agents: A timeless ramsey approach. Working Paper.

Dávila, E. & Schaab, A. (2025). Welfare assessments with heterogeneous individuals. *Journal of Political Economy*, *133*(9), 2918–2961.

De Groot, O., Mazelis, F., Motto, R., & Ristiniemi, A. (2021). A toolkit for computing constrained optimal policy projections (copps). ECB Working Paper.

Fagereng, A., Holm, M. B., & Natvik, G. J. (2021). MPC heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics*, *13*(4), 1–54.

Friedman, M. (1968). The role of monetary policy. *The American Economic Review*, *58*(1).

Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gertler, M. & Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, *7*(1), 44–76.

Giannoni, M. P. & Woodford, M. (2002). Optimal interest-rate rules. NBER Working Paper.

Greenwald, D. L., Leombroni, M., Lustig, H., & Van Nieuwerburgh, S. (2021). Financial and total wealth inequality with declining interest rates. National Bureau of Economic Research Working Paper.

Greenwood, R., Hanson, S., & Stein, J. C. (2010). A gap-filling theory of corporate debt maturity choice. *The Journal of Finance*, *65*(3), 993–1028.

Guedes, J. & Opler, T. (1996). The determinants of the maturity of corporate debt issues. *the Journal of Finance*, *51*(5), 1809–1833.

Guvenen, F., McKay, A., & Ryan, C. (2022). A tractable income process for business cycle analysis. Working Paper.

Guvenen, F., Ozkan, S., & Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, *122*(3), 621–660.

Hagedorn, M., Manovskii, I., & Mitman, K. (2019). The fiscal multiplier. National Bureau of Economic Research working paper.

Heathcote, J. & Tsujiyama, H. (2021). Optimal income taxation: Mirrlees meets ramsey. *Journal of Political Economy*, *129*(11), 3141–3184.

Hebden, J. & Winkler, F. (2024). Computation of policy counterfactuals in sequence space. Working Paper.

Holm, M. B., Paul, P., & Tischbirek, A. (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy*, *129*(10), 2861–2904.

Hoynes, H. (2000). The employment, earnings, and income of less skilled workers over the business cycle. In B. Rebecca & D. Card (Eds.), *Finding Jobs: Work and Welfare Reform* (pp. 23–71). Russell Sage Foundation.

Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, *108*(3), 697–743.

Le Grand, F., Martin-Baillon, A., & Ragot, X. (2025). Should monetary policy care about redistribution? optimal fiscal and monetary policy with heterogeneous agents. Forthcoming, Review of Economic Studies.

McKay, A. & Wolf, C. K. (2023a). Monetary policy and inequality. *Journal of Economic Perspectives*, *37*(1), 121–144.

McKay, A. & Wolf, C. K. (2023b). What can time-series regressions tell us about policy counterfactuals? *Econometrica*, *91*(5), 1695–1725.

Nuno, G. & Thomas, C. (2022). Optimal redistributive inflation. *Annals of Economics and Statistics*, (146), 3–64.

Okun, A. M. (1973). Upward mobility in a high-pressure economy. *Brookings Papers on Economic Activity*, *4*(1), 207–262.

Parker, J. A., Souleles, N. S., Johnson, D. S., & McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American economic review*, *103*(6), 2530–2553.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, *113*(4), 982–1012.

Smirnov, D. (2023). Optimal monetary policy in HANK. Universitat Pompeu Fabra working paper.

Svensson, L. E. (1997). Inflation forecast targeting: Implementing and monitoring inflation targets. *European economic review*, *41*(6), 1111–1146.

Svensson, L. E. (2005). Monetary policy with judgment: Forecast targeting.

Werning, I. (2015). Incomplete markets and aggregate demand. Working Paper 21448, National Bureau of Economic Research.

Wolf, C. K. (2023). The missing intercept: A demand equivalence approach. *American Economic Review*, *113*(8), 2232–2269.

Wolf, C. K. (2025). Interest rate cuts versus stimulus payments: An equivalence result. *Journal of Political Economy*, *133*(4), 1235–1275.

Woodford, M. (2003). *Interest and prices: Foundations of a theory of monetary policy.* princeton university press.

Yang, Y. (2024). Redistributive inflation and optimal monetary policy. University of Zurich working paper.