Optimal Policy Rules in HANK[†]

Alisdair McKay
FRB Minneapolis

Christian K. Wolf
MIT & NBER

July 17, 2022

Abstract: We characterize optimal policy rules in business-cycle models with nominal rigidities and heterogeneous households. The derived rules are expressed in terms of the causal effects of policy instruments on policymaker targets. Our first result is that the optimal policy rule of a "dual mandate" central banker—a policymaker that only cares about inflation and output—is unaffected by household heterogeneity. The optimal rule of a Ramsey planner contains an additional distributional term that incorporates the effects of her available policy instruments on consumption inequality. When calibrated to match empirical evidence on the distributional effects of monetary policy, our model implies that this concern for inequality only has a moderate effect on optimal interest rate policy. Fiscal stimulus payments, on the other hand, are strongly progressive and thus well-suited to cushion the distributional effects of cyclical fluctuations.

[†]Email: alisdair.mckay@mpls.frb.org and ckwolf@mit.edu. We thank Anmol Bhandari for helpful discussions. 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.

1 Introduction

Should household inequality affect the conduct of cyclical stabilization policy? In principle it may do so in two separate ways. First, household inequality could alter the transmission from policy instruments to any given policy target (e.g., inflation and output). This change in transmission may affect whether or not a policymaker can attain those given targets, and how policy instruments need to be set to do so. Second, household heterogeneity may also alter the policymaker targets themselves. For example, with market incompleteness, policymakers may want to dampen the distributional effects of cyclical fluctuations.

In this paper, we provide new insights to this question by expressing optimal policy rules in heterogeneous-household models as a function of empirically measurable statistics. Our first contribution is to cast optimal policy problems in a general 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, providing a general characterization of optimal policy independently of the shocks hitting the economy (Giannoni & Woodford, 2002). Drawing on McKay & Wolf (2021), we furthermore show that these optimal policy rules can be expressed in terms of the measurable causal effects of policy instruments on policymaker targets—a connection that we leverage when disciplining our quantitative analysis. With this apparatus in place, we are able to cleanly decompose the effects of inequality on stabilization policy design into its two parts: through a) policy instrument propagation and b) policymaker objectives.

Our analysis is set in a business-cycle model with nominal rigidities and household heterogeneity. Households face idiosyncratic income risk and self-insure by borrowing and saving in a (long-duration) asset, allowed to be indexed to inflation and the level of economic activity. Labor supply is intermediated by unions subject to nominal wage rigidities. The policymaker sets short-term nominal interest rates, pays transfers to households, and finances its expenditure through taxation and bond issuance.

We begin our analysis by studying the optimal policy of a conventional "dual-mandate" central banker that seeks to close the output gap and stabilize inflation. In the context of our model, this is of course an *ad hoc* loss function. We nevertheless find it useful to discuss this case, for two reasons. First, it allows us to explore the role of changes in policy instrument propagation while fixing policy targets. Second, it is arguably the relevant objective function for real-world central banks. Our main result is that the optimal interest rate target criterion is *exactly the same* as in a standard representative-agent environment. The logic is as follows. Household heterogeneity only affects the model's demand side (i.e., the "IS" curve). In the optimal policy problem, however, this demand side is a slack constraint: the policymaker can pick an output-inflation allocation subject to the Phillips curve, and then simply set nominal interest rates as necessary to generate aggregate demand consistent with the desired allocation. Household heterogeneity thus may matter for the path of the policy instrument, but not for optimal output and inflation outcomes.

We connect these theoretical results on optimal dual-mandate policy to empirical evidence on monetary policy shock transmission. The empirical literature has identified the causal effects of nominal rate changes on output and inflation (e.g. see Ramey, 2016). While silent on transmission mechanisms and so in particular on the importance of heterogeneity-related channels, aggregate data *are* informative about the size of interest rate movements required to move output and inflation by any given amount. But it then follows that *quantitatively relevant* HANK and RANK models—that is, models that are consistent with the empirical evidence on monetary shock propagation—will not only share the same optimal dual-mandate policy rule, but they will in fact also tend to agree quite closely on the interest rate path needed to implement the optimal inflation and output outcomes.[1]

We then turn to a Ramsey problem in which the planner seeks to maximize a weighted sum of individual household utilities, allowing us to explore the role of changes in policymaker objectives. Applying a second order approximation, we derive a loss function that adds to the usual output gap and inflation objectives a novel third term reflecting distributional concerns. Solving the optimal policy problem, we find that the optimal rule for any given policy instrument now contains three terms, trading off the policymaker's ability to stabilize output, inflation, and the consumption distribution. If, for example, interest rate movements do not affect consumption shares (e.g. as in Werning, 2015), then the Ramsey rule is in fact identical to the optimal dual-mandate rule; if, on the other hand, monetary policy has large redistributional effects (e.g. as in Bhandari et al., 2021), then distributional concerns may swamp price and output stability considerations.

To characterize the *empirically relevant* Ramsey policy, we again leverage the tight connection between our theoretical formulas and empirical evidence on policy shock transmission. The key new ingredient to the Ramsey rule are the causal effects of the policy instru-

[1]Formally, if estimated to be consistent with the same evidence on contemporaneous and news policy shocks, then HANK and RANK would also agree *exactly* on instrument paths (McKay & Wolf, 2021). In practice, since model estimation typically relies only on a single policy shock, agreement will be approximate.

ment on the consumption distribution. For monetary policy, we argue that the weight of the empirical evidence points to U-shaped effects on income in the cross section: households with low income and low wealth gain from expansionary policy due to a tighter labor market, while high-income and high-wealth households gain due to large capital gains. We calibrate our structural model to be consistent with these channels, and then find that it also produces a moderately pronounced U-shape on consumption, consistent with results from Norwegian administrative data reported by Holm et al. (2021).[2] Given this distributional incidence of monetary policy, we find that it is ill-suited as a tool to deal with business-cycle shocks that largely affect the bottom of the income distribution. In particular, if interest rates were set to stabilize consumption of the poor, then consumption of the rich and so aggregate output and inflation would overshoot significantly. Optimal Ramsey policy in response to such distributional shocks thus looks very similar to optimal dual-mandate policy. Those quantitative findings contrast with much recent work on optimal monetary policy with heterogeneous households that tends to find a strong role for distributional considerations (e.g. Bhandari et al., 2021; Acharya et al., 2020; Dávila & Schaab, 2022).

Finally, we turn attention to a Ramsey policymaker that jointly sets nominal interest rates and stimulus checks, consistent with recent U.S. policy practice. While equivalent in their effects on output and inflation (Wolf, 2021), the two instruments differ significantly in their distributional incidence, with fiscal stimulus payments sharply compressing consumption inequality. We thus find that the two instruments are highly complementary, with stimulus payments well-suited to offset the distributional impacts of a cyclical shock that mainly affects the bottom of the income distribution.

Literature. This paper contributes to a fast-growing literature on optimal policy design in business-cycle models with rich microeconomic heterogeneity. Conceptually, our key contribution is to characterize optimal policy through forecast target criteria, and then tie these target criteria to empirical evidence. Our computation of these optimal policy rules heavily leverages sequence-space representations of equilibria and thus the recent work of Auclert et al. (2021).[3] A contemporaneous paper that does the same is Dávila & Schaab

[2]Capital gains—which largely accrue to advantaged groups—are large in dollar terms, but their immediate pass-through to consumption is relatively low. The labor income gains that flow to low-income, high-MPC households, on the other hand, pass through immediately to consumption. The overall effects of interest rate policy on the consumption distribution are thus much more equitable than the effects on household balance sheets. Our model matches these empirical patterns.

[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-

(2022). Those authors do not rely on linear-quadratic approximations, thus providing more general optimal policy results, but without our tight connection to empirical evidence. Our discussion of the mapping between policy rules and empirical evidence builds on our earlier work in McKay & Wolf (2021). In operationalizing these results, we are fortunate to rely on a recent literature that empirically documents the distributional effects of monetary policy interventions (e.g. Holm et al., 2021; Andersen et al., 2021; Bartscher et al., 2021).

Our conceptual innovations allow us to revisit several substantive results in the prior heterogeneous-household optimal policy literature. First, we derive a sharp set of irrelevance results for the effects of household inequality on optimal policy design. Prior work has emphasized that household inequality will affect policy propagation by altering the demand side of the economy (Kaplan et al., 2018; Auclert, 2019); our analysis however reveals that these changes leave the optimal target criterion and thus equilibrium paths of inflation and output unaffected. If furthermore matched to the same empirical evidence on policy shock propagation, then heterogeneous-agent and representative-agent models will also (roughly) agree on the required interest rate paths. Second, cyclical stabilization policy may be used to provide insurance (Acharya et al., 2020; Bhandari et al., 2021; Le Grand et al., 2021). We contribute to this area by characterizing optimal policy in a model that matches empirical evidence on the distributional consequences of stabilization policy and thus on its scope for insurance. Third, our analysis of optimal joint fiscal-monetary policy extends results in Wolf (2021) and Bilbiie et al. (2021). Wolf (2021) shows that, in a standard HANK environment, transfers and monetary policy can implement the same set of aggregate allocations, but differ in their distributional implications. These are purely positive properties of the model. On the normative side, in a two-agent environment, Bilbiie et al. (2021) argue that monetary policy and fiscal policy can be used together to stabilize both the aggregate level of activity and the consumption shares of the two household types. Our work is complementary: we analyze the optimal monetary-fiscal policy mix in an environment with rich heterogeneity and a tight link to empirical evidence.[4]

space linear-quadratic policy problems have been used in prior work to derive "optimal policy projections" (Svensson, 2005; De Groot et al., 2021; Hebden & Winker, 2021). We show how to fit optimal policy problems in HANK models into this environment.

[4]A fourth theme of the recent optimal policy literature is that inequality among households introduces a new source of time inconsistency (Acharya et al., 2020; Nuno & Thomas, 2021; Dávila & Schaab, 2022). 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.

Outline. In Section 2, we present a general linear-quadratic sequence-space problem and its solution. Section 3 then outlines our HANK model, and Sections 4 and 5 present our results for optimal dual-mandate and Ramsey policy, respectively. We conclude in Section 6.

2 Linear-quadratic problems in the sequence space

Throughout this paper we study optimal policy problems that can be recast as deterministic linear-quadratic control problems. We begin in Section 2.1 by stating the problem and presenting its solution. Section 2.2 then places our analysis in the broader context of the literature, focussing in particular on the tight connection between our expressions and empirical evidence on the effects of policy shocks.

2.1 Problem & solution

We consider a policymaker that faces a linear-quadratic optimal control problem.

Preferences. The policymaker targets I variables, indexed by i. We let x_{it} be the deviation of the ith target variable from its target value at date t. We then consider a policymaker with quadratic loss function

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

where $\boldsymbol{x}_i \equiv (x_{i0}, x_{i1}, \dots)'$ is the perfect-foresight sequence of the ith target variable through time and $\boldsymbol{x} \equiv (\boldsymbol{x}_1, \boldsymbol{x}_2, \dots, \boldsymbol{x}_I)'$ stacks those paths for all of the I targets. The λ_i's denote the weights associated with the different policy targets, with $\Lambda \equiv \mathrm{diag}(\lambda_1, \lambda_2, \dots, \lambda_I)$. Finally $W = \mathrm{diag}(1, \beta, \beta^2, \dots)$ summarizes the effects of discounting in the policymaker preferences, with discount factor $\beta \in (0, 1)$.

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

$$\mathcal{H}_x\boldsymbol{x} + \mathcal{H}_z\boldsymbol{z} + \mathcal{H}_\varepsilon\boldsymbol{\varepsilon} \;=\; \boldsymbol{0}, \tag{2}$$

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.

While the structural models considered in the remainder of this paper directly map into constraints of the general form (2), it follows from the discussion in McKay & Wolf (2021) that these constraints can equivalently and more conveniently be expressed as

$$\boldsymbol{x}_i = \sum_{j=1}^{J} \Theta_{x_i,z_j} \boldsymbol{z}_j + \sum_{q=1}^{Q} \Theta_{x_i,\varepsilon_q} \boldsymbol{\varepsilon}_q, \quad i = 1, 2, \ldots, I \tag{3}$$

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 (3) thus expresses the policy targets directly in terms of impulse responses to policy instruments and exogenous shocks, as opposed to imposing implicit relationships as in (2).[5]

Problem & solution. The optimal policy problem is to choose the instrument paths $\boldsymbol{z}$ to minimize (1) subject either to (2) (for the original constraint formulation) or (3) (for the simplified re-cast constraint). The policymaker thus minimizes a convex objective subject to linear constraints, and so the first-order conditions are necessary and sufficient for a solution to the problem.

For intuition, it is simpler and more instructive to use the constraint (3). Minimizing (1) subject to (3) yields:

1. **Optimal policy rule.** For each policy instrument $\boldsymbol{z}_j$, the paths of the policy targets satisfy the "policy criterion"

$$\sum_{i=1}^{I} \lambda_i \cdot \underbrace{\Theta'_{x_i,z_j} W}_{\text{(discounted) causal effect of } z_j \text{ on } x_i} \cdot \boldsymbol{x}_i = \mathbf{0}, \quad j = 1, 2, \ldots, J \tag{4}$$

   (4) is simply the first-order condition of the optimal policy problem. It says that, for each instrument $\boldsymbol{z}_j$, the paths of the policy 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

[5] The equivalence of (2) and (3) would be immediate for invertible $\mathcal{H}_x$. In typical macroeconomic models, however, $\mathcal{H}_x$ is not invertible, so recasting the constraint as (3) requires additional arguments. McKay & Wolf (2021) 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.

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 emphasize two important features of such rules. First, they are derived without reference to and so apply independently of the non-policy shocks hitting the economy. This robustness property is one of the main virtues of target policy criteria (Giannoni & Woodford, 2002). Second, note that the optimal policy rule for instrument j places no weight on a policy target i that cannot be moved by instrument j (i.e., $\Theta_{x_i,z_j} = \mathbf{0}$), even if $\lambda_i > 0$—intuitively, if an instrument cannot affect a target, then this target plays no role in informing the setting of the instrument.

2. **Optimal policy path.** Given the exogenous shock paths $\boldsymbol{\varepsilon}$, the policy rule (4) together with the constraints (3) characterizes the evolution of the dynamic system. In particular, the optimal instrument path $\boldsymbol{z}^*$ satisfies

$$\boldsymbol{z}^* \;\equiv\; -\left(\Theta'_{x,z}(\Lambda \otimes W)\Theta_{x,z}\right)^{-1} \times \left(\Theta'_{x,z}(\Lambda \otimes W)\Theta_{x,\varepsilon}\,\boldsymbol{\varepsilon}\right), \tag{5}$$

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. The optimal path of the policy instruments thus has an intuitive regression interpretation: the instruments $\boldsymbol{z}$ are set to offset as well as possible—in a weighted least-squares sense—the perturbation to the policy targets $\boldsymbol{x}$ caused by the exogenous shocks, given as $\Theta_{x,\varepsilon} \times \boldsymbol{\varepsilon}$. In particular, the policymaker will rely most heavily on the tools z_j that are best suited to offset the perturbation to its targets induced by a particular shock path $\boldsymbol{\varepsilon}$.

2.2 Discussion

Equations (4) and (5) will guide our analysis in much of the remainder of the paper. In this section we briefly relate our results to prior work on: first, stochastic linear-quadratic problems; and second, empirical measurement of the propagation of policy shocks.

Deterministic transitions vs. aggregate risk. It is well-established that, by certainty equivalence, the first-order perturbation solution of models with aggregate risk is mathematically identical to linearized perfect-foresight transition paths.[6] This insight im-

[6]For detailed discussions of this point see for example Fernández-Villaverde et al. (2016), Boppart et al. (2018) or Auclert et al. (2021).

plies the following connections between our linear-quadratic perfect foresight problem and the canonical linear-quadratic stochastic problem (as in Benigno & Woodford, 2012). First, the policy target criterion (4) corresponds to a *forecast* targeting criterion in a stochastic economy. For a time-0 problem with commitment, that forecast targeting criterion is simply

$$\mathbb{E}_0 \left[\sum_{i=1}^{I} \lambda_i \cdot \Theta'_{x_i, z_j} W \cdot \boldsymbol{x}_i \right] = \mathbf{0}, \quad j = 1, 2, \ldots, J \tag{6}$$

This is an implicit rule that determines the expected evolution of the economy as of date 0. In a stochastic environment, new shocks will occur as time goes by, causing the evolution of the economy to deviate from what was expected at date 0. In this case, (6) gives a rule for how to *revise* forecasts at each date. Second, by the same logic, the optimal instrument path $\boldsymbol{z}^*$ in (5) corresponds to the instrument *impulse response* to a time-0 shock that changes expectations of the exogenous shifters from $\mathbf{0}$ to $\boldsymbol{\varepsilon}$. The exact same impulse response interpretation applies to our solution for the paths of the policy targets $\boldsymbol{x}^*$.

Measurement. The linear maps stacked in $\Theta_{x,z}$ collect the dynamic causal effects of variations in the policy instruments z onto the policymaker targets x. As we show formally in McKay & Wolf (2021), entries of these maps can be estimated using semi-structural time series methods applied to identified policy *shocks* (as in e.g. Ramey, 2016). We will leverage this connection in our quantitative analysis in Sections 4.3 and 5.4.

Outlook. In the remainder of this paper we will first show that optimal policy problems in models with household heterogeneity can be represented in the form of our linear-quadratic control problem, and then use (4) and (5) to characterize optimal policy rules, throughout connecting as closely as possible to empirical evidence on policy shock propagation. Section 3 begins the analysis with a description of our model environment.

3 Model

Our model environment is a standard HANK economy, with two somewhat special features. First, our model features sticky wages. While the early HANK literature focussed on sticky price models (McKay et al., 2016; Kaplan et al., 2018), some recent contributions have shifted their focus to frictions in wage-setting. Appealingly, such frictions generate more realistic responses of capital income to changes in aggregate demand (Broer et al., 2020); furthermore, they allow the model to be consistent with small household marginal propensities to earn (see Auclert et al., 2020). Second, households in our model save in a long-duration asset. Including a long-duration asset allows us to more flexibly capture the consequences of policy actions for the valuation of assets and liabilities. Both of these changes will be important to connect our model to empirical evidence on policy propagation.

Time is discrete and runs forever, $t = 0, 1, 2, \ldots$. Consistent with our linear-quadratic framework in Section 2, we will consider linearized perfect-foresight transition sequences. By certainty equivalence, our solutions will be identical to the analogous economy with aggregate risk, solved using conventional first-order perturbation techniques with respect to aggregate variables. Throughout this section, boldface denotes time paths (so e.g., $\boldsymbol{x} \equiv (x_0, x_1, x_2, \ldots)'$), bars indicate the model's deterministic steady state ($\bar{x}$), and hats denote (log-)deviations from the steady state ($\widehat{x}$).[7]

3.1 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

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{it}^{1-\gamma} - 1}{1-\gamma} - \nu\left(\ell_{it}\right) \right], \tag{7}$$

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

Households are endowed with stochastic idiosyncratic labor productivity e_{it}. We let ζ_{it} be a stochastic event that determines the labor productivity of household i at date t. We then assume there is a function Φ that maps ζ_{it} to e_{it},

$$e_{it} = \Phi(\zeta_{it}, m_t, y_t).$$

This mapping potentially depends on an exogenous distributional shock, m_t, and an endogenous component captured by aggregate income, y_t. ζ_{it} itself follows a stationary Markov process. A canonical heterogeneous agent model would simply set $e_{it} = \zeta_{it}$. We further assume that $\int e_{it} di = 1$ for any value of m_t and y_t, so these variables only affect the distribution of labor productivities (and not the average level). For the quantitative analysis in

[7] To be precise, we use log deviations for the variables $\{y, c, \ell, 1+r, 1+i, v\}$ and level deviations for the variables $\{\pi, \tau_x, \tau_e, m\}$.

Sections 4.3 and 5.4, the shock m_t will be our example of an inequality shock—a shock that affects aggregate demand through redistribution and precautionary savings motives.

Total pre-tax nominal household labor income is $e_{it}w_t\ell_{it}$, where w_t is an aggregate nominal wage rate per efficiency unit of labor. As we describe below, labor supply is determined by a labor market union, so hours worked ℓ_{it} are taken as given by the household. Total labor income is taxed at some constant proportional rate τ_y. Households also 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 refer to it as a "fiscal 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. Finally, households can borrow and save through a financial asset with realized time-t real return r_t, subject to an exogenous borrowing constraint $\underline{a} \leq 0$. Putting all the pieces together, the budget constraint is

$$a_{it} + c_{it} = (1 + r_t)a_{it-1} + (1 - \tau_y)e_{it}\frac{w_t}{p_t}\ell_{it} + \tau_{x,t} + \tau_{e,t}e_{it}, \tag{8}$$

where a_{it} are assets held at the end of period t and p_t is the nominal price of the final good. While asset returns are expressed here in real terms, the contracts that give rise to these returns can be nominal, as discussed further below.

The solution to each individual household i's consumption-savings problem gives a mapping from paths of wages $\boldsymbol{w}$, hours worked $\boldsymbol{\ell}_i$, real returns $\boldsymbol{r}$, transfers $\boldsymbol{\tau}_x$ and $\boldsymbol{\tau}_e$, prices $\boldsymbol{p}$ and shocks $\boldsymbol{m}$ to that household's consumption $\boldsymbol{c}_i$. Aggregating consumption decisions across all households, we thus obtain an aggregate consumption function $\mathcal{C}(\bullet)$, exactly as in Auclert et al. (2018) or Wolf (2021):

$$\boldsymbol{c} = \mathcal{C}(\boldsymbol{w}/\boldsymbol{p}, \boldsymbol{\ell}, \boldsymbol{r}, \boldsymbol{\tau}_x, \boldsymbol{\tau}_e, \boldsymbol{m}), \tag{9}$$

where $\boldsymbol{w}/\boldsymbol{p}$ is the sequence of real wages. Linearizing this consumption function around the deterministic steady state yields

$$\widehat{\boldsymbol{c}} = \mathcal{C}_{w/p}\widehat{\boldsymbol{w}/\boldsymbol{p}} + \mathcal{C}_\ell\widehat{\boldsymbol{\ell}} + \mathcal{C}_r\widehat{\boldsymbol{r}} + \mathcal{C}_x\widehat{\boldsymbol{\tau}}_x + \mathcal{C}_e\widehat{\boldsymbol{\tau}}_e + \mathcal{C}_m\widehat{\boldsymbol{m}}, \tag{10}$$

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

3.2 Technology, unions, and firms

Labor supply is intermediated by a unit continuum of labor unions, and a competitive producer then packages union labor supply to produce the final good. Since this production model block is relatively standard, we only state and briefly discuss the key relations here, with a detailed discussion relegated to Appendix A.1.

Union k demands ℓ_{ikt} units of labor from household i. The final good is sold at nominal price p_t and produced by aggregating the labor supply of all individual unions k, denoted $\ell_{kt} \equiv \int_0^1 e_{it}\ell_{ikt}di$. The aggregate production function takes a standard constant elasticity form, with the elasticity of substitution between labor of different unions, η_t, allowed to vary exogenously over time. This "supply"-type shock will be important in our discussion of optimal policy: it implies changes in market power that result in inefficient fluctuations in the flexible-price level of output, thus creating a trade-off between stabilizing output around its efficient level and stabilizing inflation. All unions satisfy labor demand by rationing labor equally across all households. This rationing rule together with marginal cost pricing ($w_t = p_t$) for the competitive producer imply that $e_{it}\ell_{it}\frac{w_t}{p_t} = e_{it}y_t$ for all i.

Each union sets its nominal wage in standard Calvo fashion, with probability $1-\theta$ of updating the wage each period. As usual, unions select their wages upon reset based on current and future marginal rates of substitution between leisure and consumption among its household members. Given separable preferences and with everyone supplying an equal amount of hours worked, it follows that all households share a common marginal disutility of labor. The marginal utility of consumption, however, is generally not equalized. For reasons that we will discuss in detail later, we assume that the union evaluates the benefits of higher after-tax income using the marginal utility of average consumption ($c_t^{-\gamma}$) rather than the average of marginal utilities ($\int_0^1 c_{it}^{-\gamma}di$), as also done in Wolf (2021) and Auclert et al. (2021). We show in Appendix A.1 that the solution to this union problem then gives rise to a standard linearized perfect-foresight New Keynesian Phillips curve:

$$\widehat{\pi}_t = \kappa\widehat{y}_t + \beta\widehat{\pi}_{t+1} + \psi\widehat{\eta}_t, \tag{11}$$

where $\kappa \equiv (\phi+\gamma)\frac{(1-\theta)(1-\beta\theta)}{\theta}$, $\phi \equiv \frac{\nu_{\ell\ell}(\bar{\ell})\bar{\ell}}{\nu_\ell(\bar{\ell})}$ and $\psi \equiv -\frac{\kappa}{(\phi+\gamma)(\bar{\eta}-1)}$. We allow for a (time-invariant) subsidy on union labor hiring, financed with lump-sum taxes also levied on the unions; this subsidy will matter in Section 5, where we require efficiency of the deterministic steady state to write our optimal policy problem in a form consistent with the linear-quadratic set-up of Section 2, exactly as in prior work (e.g. Woodford, 2003).

Finally, aggregate production is equal to

$$y_t = \frac{\ell_t}{d_t}, \tag{12}$$

where $\ell_t \equiv \int_0^1 \int_0^1 e_{it}\ell_{ikt} didk$ and $d_t \geq 1$ captures efficiency losses related to aggregate wage dispersion across unions.

3.3 Asset structure

There are two different assets in the economy: a short-term, risk-free nominal bond that is in zero net supply, and a second asset that is long-lived, (partially) indexed to inflation and output, and in positive net supply. By arbitrage, both assets will provide the same expected returns along equilibrium transition paths, thus allowing us to consider a single asset in the household budget constraint (8). The realized return at date 0, however, will differ between the two assets. The purpose of the long-term asset is to allow monetary policy to have data-consistent effects on household asset income including capital gains—a key determinant of the policy's distributional implications and so, as we will see, optimal policy design.

A unit of the nominal bond purchased at time t returns $\frac{1+i_t}{1+\pi_{t+1}}$ units of the final good at time $t+1$. For the second asset, at time t, households can purchase a unit of the asset for a real price of q_t (i.e., denominated in goods); at time $t+1$, the household receives a real "coupon" of $(\bar{r}+\delta)(1+\pi_{t+1})^{\chi_\pi - 1}(\frac{y_{t+1}}{\bar{y}})^{\chi_y}$ and furthermore retains a fraction $(1-\delta)(1+\pi_{t+1})^{\chi_\pi-1}$ of the asset position, now valued at $(1-\delta)(1+\pi_{t+1})^{\chi_\pi-1}q_{t+1}$ in units of goods. This asset set-up captures the following features. First, the parameter δ controls the maturity of the asset, with coupons decaying at rate δ. The coupon scaling factor $(\bar{r}+\delta)$ normalizes the steady-state price of the bond to one. Second, the inflation term captures inflation indexation, with $\chi_\pi = 1$ corresponding to a real bond, $\chi_\pi = 0$ corresponding to a nominal bond, and $\chi_\pi \in (0,1)$ giving the intermediate case. Third, the output term captures the sensitivity of asset income with respect to real economic activity, with $\chi_y > 0$ corresponding to an asset with higher returns in good times.

Now let r_t denote the real return on the second asset. Bond coupon payments are then linked to real returns as

$$1 + r_t = \frac{(\bar{r}+\delta)(1+\pi_t)^{\chi_\pi-1}\left(\frac{y_t}{\bar{y}}\right)^{\chi_y} + (1-\delta)(1+\pi_t)^{\chi_\pi-1}q_t}{q_{t-1}}, \tag{13}$$

with the convention that $q_{-1} = \bar{q}$. (13) gives real bond returns as a function of bond prices, inflation, and real output. By arbitrage, for all $t > 0$, real returns are furthermore linked to returns on the nominal bond via the standard Fisher relation

$$1 + r_t = \frac{1 + i_{t-1}}{1 + \pi_t}. \tag{14}$$

At date $t = 0$, the realized return on a household's portfolio will depend on the composition of its portfolio between the two assets. We assume that there are no existing gross positions in the short-term asset, so time-0 realized returns are simply those on the long-term asset, for all households.

3.4 Government

The final actor in our model is the government. The government collects tax revenue, pays out lump-sum transfers, sets the nominal rate on the short-term bond, and issues positive quantities of the long-lived asset. Letting b_{t-1} denote outstanding claims on the government (denominated in units of bonds), the government budget constraint becomes

$$(\bar{r} + \delta)(1 + \pi_t)^{\chi_\pi - 1} \left(\frac{y_t}{\bar{y}}\right)^{\chi_y} b_{t-1} + \tau_{x,t} + \tau_{e,t} = \tau_y y_t + q_t \left(b_t - (1 - \delta)(1 + \pi_t)^{\chi_\pi - 1} b_{t-1}\right). \tag{15}$$

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:

$$\tau_{e,t} = (\bar{r} + \sigma)(b_{t-1} - \bar{b}). \tag{16}$$

When the stock of bonds outstanding exceeds its steady state level, taxes are raised to pay interest and a portion σ of the outstanding bonds. For example, if $\sigma = \delta$, then the policymaker pays off interest and any amount of maturing bonds beyond the steady state level. Given this feedback rule, government debt then evolves according to (15).

3.5 Equilibrium

We can now define a linearized perfect-foresight transition equilibrium in this economy.[8]

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 $\{i_t, \tau_{x,t}, \tau_{e,t}, b_t\}_{t=0}^{\infty}$ and a set of aggregates $\{c_t, y_t, a_t, \pi_t, r_t, q_t, w_t/p_t, \ell_t\}_{t=0}^{\infty}$ 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 $\{a_t\}_{t=0}^{\infty}$ is consistent with the budget constraint (8), aggregated across households.*
2. *The real wage is consistent with marginal cost pricing for final goods firms, so $w_t = p_t$.*
3. *The paths of $\{\ell_t, y_t\}_{t=0}^{\infty}$ satisfy the aggregate production function (12).*[9]
4. *The paths $\{\pi_t, y_t, \eta_t\}_{t=0}^{\infty}$ are consistent with the Phillips curve (11).*
5. *The evolution of government debt b_t and the endogenous component of transfers $\tau_{e,t}$ are consistent with the budget constraint (15) and law of motion (16).*
6. *The asset returns $\{r_t, i_t, q_t\}_{t=0}^{\infty}$ satisfy (13) and (14).*
7. *The output and asset markets clear, so $y_t = c_t$ and $a_t = q_t b_t$.*

Sections 4 and 5 will describe optimal policy problems and so discuss how the policy instrument paths $\{i_t, \tau_{x,t}\}$—simply taken as given in Definition 1—are determined.

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*

$$\widehat{\boldsymbol{\pi}} \quad = \quad \kappa \widehat{\boldsymbol{y}} + \beta \widehat{\boldsymbol{\pi}}_{+1} + \psi \widehat{\boldsymbol{\eta}}, \tag{17}$$

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

[9] 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 social welfare function in Section 5.

$$\widehat{\boldsymbol{y}} \;=\; \tilde{\mathcal{C}}_y\widehat{\boldsymbol{y}} + \tilde{\mathcal{C}}_\pi\widehat{\boldsymbol{\pi}} + \tilde{\mathcal{C}}_i\widehat{\boldsymbol{i}} + \tilde{\mathcal{C}}_x\widehat{\boldsymbol{\tau}}_x + \mathcal{C}_m\widehat{\boldsymbol{m}}, \tag{18}$$

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 in Definition 1 to two equations: the Phillips curve (17) (which is simply a stacked perfect-foresight version of the original relation (11)); and the IS curve (18), which differs from the consumption function (10) chiefly in that it imposes: (i) output market-clearing ($\widehat{y}_t = \widehat{c}_t$); (ii) the aggregate production function ($\widehat{\ell}_t = \widehat{y}_t$); (iii) the equilibrium real wage ($w_t/p_t = 1$); and (iv) 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)? *Positively*, the main change is that a simple aggregate Euler equation,

$$\widehat{y}_t = -\frac{1}{\gamma}(\widehat{i}_t - \widehat{\pi}_{t+1}) + \widehat{y}_{t+1}, \tag{19}$$

is now replaced by a more general IS curve (18). 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 (17)—kept exactly as in standard representative-agent models. To arrive at this clean separation, our assumptions on union bargaining (see Section 3.2) are central. We adopt this approach because the demand-side effects of household heterogeneity are the focus of the recent HANK literature (Kaplan et al., 2018; Auclert et al., 2018).[10] *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. First, in Section 4, we isolate the role of changes in propagation by studying a dual-mandate optimal policy problem. Then, in Section 5, we turn to the full Ramsey problem, thus allowing the planner objective to change.

[10]Furthermore, as discussed in Wolf (2021), the supply-side effects of household heterogeneity are generally of limited importance, because (i) the dynamics of average marginal utilities and marginal utilities at the average rarely differ much, and (ii) the importance of any remaining difference is dampened by wage rigidity.

4 Optimal dual-mandate policy

In this section we study the optimal policy problem of a conventional dual-mandate policymaker; that is, a policymaker that simply seeks to stabilize fluctuations in aggregate inflation and the output gap. In the context of the structural model of Section 3, such a loss function is *ad hoc*, but we find it interesting, for two reasons. First, it is conceptually useful, as it allows us to isolate the role of inequality for optimal policy design through its effects on policy propagation. Second, it is practically relevant, as real-world central banks are often mandated to achieve these types of objectives.

We begin in Section 4.1 by stating the optimal policy problem in linear-quadratic form. We then in Section 4.2 characterize the solution in the form of a policy forecast target criterion. Finally, in Section 4.3, we present some quantitative explorations, leveraging in particular the connection between our optimal policy formulas and empirical evidence.

4.1 The optimal policy problem

We consider a policymaker with objective function

$$\mathcal{L}^{DM} \equiv \sum_{t=0}^{\infty} \beta^t \left[\lambda_\pi \widehat{\pi}_t^2 + \lambda_y \widehat{y}_t^2\right]. \tag{20}$$

(20) is a *dual-mandate* loss function: the policymaker wishes to stabilize inflation and output around the deterministic steady state, with weights λ_π and λ_y, respectively.

For most of this section, we focus on the optimal setting of nominal interest rates i_t, with only brief reference to optimal stimulus check policy.[11] The policymaker sets nominal interest rates to minimize (20) subject to the equilibrium constraints embedded in Definition 1. By Lemma 1, we can reduce these two constraints to two simple linear relationships: the Phillips curve (17) and the IS curve (18). This optimal policy problem is a minimal departure from optimal policy analysis in conventional representative-agent environments: the loss function (by assumption) and the supply side are unaffected, while the demand constraint changes from a simple aggregate Euler equation as in (19) to the richer demand relation (18).

Note that so far the constraints of this policy problem take the form of our general linear

[11] Of course, since our model does not feature Ricardian households, monetary and fiscal policy are never truly separate. Thus, to be more precise, we study the choice of nominal interest rates conditional on the debt feedback rule (16).

constraint (2). By the arguments in McKay & Wolf (2021), we can equivalently re-write this constraint set in impulse response space, thus giving our alternative formulation (3).[12] For future reference, we write the constraints in impulse response space as

$$\widehat{\boldsymbol{\pi}} = \Theta_{\pi,i}\widehat{\boldsymbol{i}} + \Theta_{\pi,x}\widehat{\boldsymbol{\tau}}_x + \Theta_{\pi,\eta}\widehat{\boldsymbol{\eta}} + \Theta_{\pi,m}\widehat{\boldsymbol{m}}, \tag{21}$$

$$\widehat{\boldsymbol{y}} = \Theta_{y,i}\widehat{\boldsymbol{i}} + \Theta_{y,x}\widehat{\boldsymbol{\tau}}_x + \Theta_{y,\eta}\widehat{\boldsymbol{\eta}} + \Theta_{y,m}\widehat{\boldsymbol{m}}. \tag{22}$$

Computational details. Solving the dual-mandate optimal policy problem is straightforward. Key to this computational simplicity is that the maps characterizing the linear-quadratic problem—either the $\tilde{\mathcal{C}}$'s in the original constraint formulation or the Θ's in the equivalent impulse response space formulation—can be obtained straightforwardly as a side-product of standard sequence-space solution output, following the methods developed in Auclert et al. (2021). It thus follows that optimal policy analysis in the dual-mandate case comes at essentially zero additional computational cost: if a researcher can solve her HANK model *given* a policy rule, then she is only a trivial linear-quadratic problem away from also obtaining an *optimal* policy rule for a given quadratic loss.

4.2 Policy rule irrelevance

Our first main result is that, under mild regularity conditions on the linear map $\tilde{\mathcal{C}}_i$—i.e., the mapping from nominal interest rate paths to net excess consumption demand in the generalized "IS" curve (18)—, the optimal monetary policy forecast target criterion is completely unaffected by household heterogeneity.

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*

$$\widehat{\boldsymbol{c}} \in \textit{image}(\tilde{\mathcal{C}}_i), \tag{23}$$

then the optimal monetary policy rule for a dual-mandate policymaker with loss function (20)

[12] In McKay & Wolf (2021), we close the model with a determinacy-inducing policy *rule* for the policy instruments, here i and τ_x. The causal effect matrices are then defined as impulse response matrices for shocks to the baseline rule. For example, if the inflation and interest rate impulse response matrices to monetary shocks to the base rule are denoted $\tilde{\Theta}_{\pi,i}$ and $\tilde{\Theta}_{i,i}$, then $\Theta_{\pi,i} \equiv \tilde{\Theta}_{\pi,i}\tilde{\Theta}_{i,i}^{-1}$. This re-writing is without loss of generality (see McKay & Wolf, 2021).

can be written as the forecast target criterion

$$\lambda_\pi \widehat{\pi}_t + \frac{\lambda_y}{\kappa}\left(\widehat{y}_t - \widehat{y}_{t-1}\right) = 0, \quad \forall t = 0, 1, \ldots \tag{24}$$

Recall that our expressions for general linear-quadratic policy problems derived in Section 2 immediately yield the optimal dual-mandate policy target criterion as

$$\lambda_\pi \cdot \Theta'_{\pi,i} \cdot \widehat{\boldsymbol{\pi}} + \lambda_y \cdot \Theta'_{y,i} \cdot \widehat{\boldsymbol{y}} = \mathbf{0}. \tag{25}$$

The proof of Proposition 1 leverages the structure of our particular model to turn the general expression (25) into the simple rule (24). Importantly, the rule (24) is *exactly* the same as in conventional representative-agent optimal policy analyses (e.g. as in Galí, 2015; Woodford, 2003). The logic underlying this result is as follows. In the familiar representative-agent policy problem, any desired path of output and inflation that is consistent with the Phillips curve can be implemented through a suitable choice of interest rates. The IS curve is thus a slack constraint: the policymaker picks the best output-inflation pair subject to the Phillips curve constraint, and then sets interest rates residually to deliver the required time path of demand. Our technical condition in (23) is precisely enough to ensure that this logic carries through in our environment with household heterogeneity. In words, 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. The proof of Proposition 1 reveals that this is sufficient to ensure that any desired output-inflation pair consistent with the Phillips curve (17) is in fact implementable. But then, with the Phillips curve as the supply side of the economy not depending on household inequality, we find that the target criterion is the same as in conventional representative-agent models.

The implementability condition (23) is discussed further in Wolf (2021). That paper shows—analytically in simple models, and numerically in heterogeneous-agent environments—that interest rate policies are indeed generally flexible enough to induce every possible zero net present value path of aggregate net excess demand. The irrelevance of household heterogeneity for optimal forecast criteria is thus a robust feature of HANK-type environments.

Implications for monetary policy practice. The upshot of Proposition 1 is that, independently of the non-policy shocks hitting the economy, under optimal dual-mandate policy, the equilibrium paths of output and inflation will be unaffected by household heterogeneity and thus equal to those in a standard representative-agent economy. The only possible effect of heterogeneity is to change the instrument paths—i.e., the current and future values of nominal interest rates—required to attain those desired output and inflation paths. We conclude that the positive implications of household heterogeneity have a rather limited effect on the practice of dual-mandate policymakers: they can continue to set their instruments to bring projections of macroeconomic outcomes in line with target, exactly as done in standard practice of flexible inflation targeting.[13]

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 (2021), who identifies conditions under which interest rate and stimulus check policies can implement the same sequences of aggregate output and inflation. More formally, it follows from his results that, if

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

for all sequences $\widehat{\boldsymbol{c}}$ with zero net present value, then stimulus check policies can also implement the target criterion (24), 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 *joint* optimal policy problem 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. Section 5 considers an alternative resolution to this indeterminacy: a richer loss function.

4.3 Quantitative analysis & connection to empirical evidence

We have seen that household heterogeneity does not affect the optimal inflation and output gap outcomes implemented by a dual-mandate policymaker. Heterogeneity could, however, in principle quite materially affect the time paths of nominal rates required to achieve those optimal outcomes. This section leverages the close connection between our theory and empirical evidence on policy shock propagation to argue that, *in quantitatively relevant models*,

[13] Bernanke (2015) succinctly summarizes the salience of this perspective for Federal Reserve policymaking practice: *"The Fed has a rule. The Fed's rule is that we will go for a 2 percent inflation rate. We will go for the natural rate of unemployment. We put equal weight on those two things. We will give you information about our projections about our interest rates. That is a rule and that is a framework that should clarify exactly what the Fed is doing."*

the effects of household heterogeneity on optimal rate paths are likely to be modest in scope.

Exact instrument path irrelevance. We begin with another exact irrelevance result, building closely on McKay & Wolf (2021). For this result, it will prove convenient to re-state (5) for the optimal policy instrument path, here specialized to the policy instrument i and for the policy targets $x = (y, \pi)'$, for some generic set of shocks $\boldsymbol{\varepsilon}$:

$$\widehat{\boldsymbol{i}}^* \equiv -\left(\Theta_{x,i}'(\Lambda \otimes W)\Theta_{x,i}\right)^{-1} \times \left(\Theta_{x,i}'(\Lambda \otimes W)\Theta_{x,\varepsilon} \cdot \boldsymbol{\varepsilon}\right). \tag{27}$$

Equation (27) has the following important implication. Any two models—say HANK and RANK—that agree on (i) the effects of a given shock $\boldsymbol{\varepsilon}$ on output and inflation, $\Theta_{y,\varepsilon} \cdot \boldsymbol{\varepsilon}$ and $\Theta_{\pi,\varepsilon} \cdot \boldsymbol{\varepsilon}$, and (ii) the effects of interest rate changes on output and inflation, $\Theta_{y,i}$ and $\Theta_{\pi,i}$, will *necessarily* agree on the optimal interest rate path $\widehat{\boldsymbol{i}}^*$.

We view this irrelevance result as informative because its ingredients are measurable. In particular, the dynamic causal effects of interest rate changes on aggregate outcomes—that is, elements of $\Theta_{y,i}$ and $\Theta_{\pi,i}$—are the estimands of a large empirical literature on monetary policy shocks (Ramey, 2016). Structural models are often calibrated or estimated to be consistent with estimates from this literature, which leads them to yield similar outcomes for $\Theta_{y,i}$ and $\Theta_{\pi,i}$. The degree to which such quantitative models can disagree on optimal dual-mandate interest rate paths is thus limited by the empirical evidence.

Quantitative illustration. We close with a quantitative illustration of the analytical results presented in this section. To this end, we study optimal dual-mandate monetary policy in response to a cost-push shock η_t in a calibrated version of our HANK model. Details of the calibration are postponed until Section 5.3; for purposes of the discussion here, it suffices to note that the model has been parameterized to in particular be consistent with empirical evidence on the output gap and inflation effects of monetary shocks. We then contrast aggregate outcomes in this economy with those in an analogous RANK model, calibrated similarly to be consistent with empirical policy shock evidence.

Results are displayed in Figure 1. We emphasize the following two takeaways, consistent with our analytical discussion in the rest of this section. First, consistent with Proposition 1, the output and inflation paths are exactly the same. The optimal trade-off between inflation and output is fully governed by the Phillips curve, which itself is not affected by household inequality. Second, the nominal interest rate path required to implement the optimal outcome is quite similar across the two models. Intuitively, since by construction both models

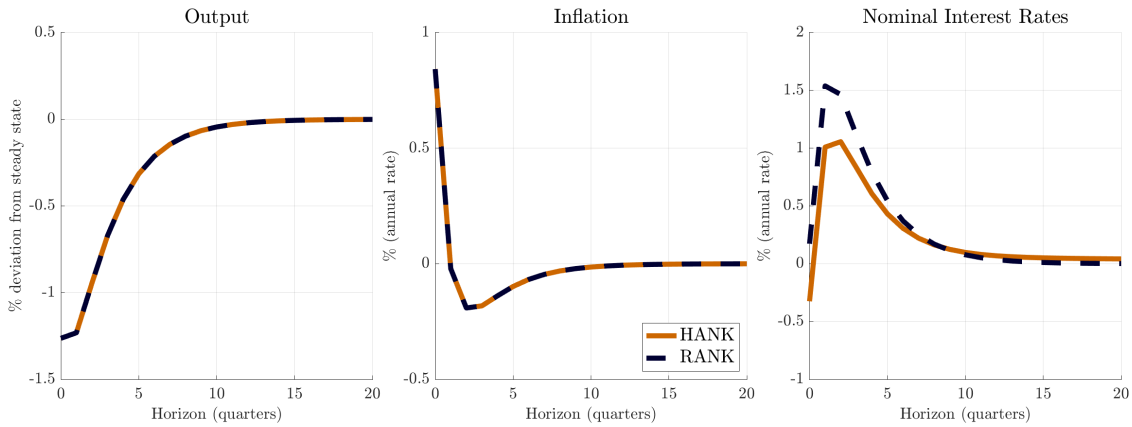

Figure 1: Optimal dual-mandate policy response to a cost-push shock. The RANK model replaces our general "IS" curve (18) with the simple textbook Euler equation (19). We then calibrate all parameters as in the headline HANK model (see Section 5.3), with one exception: we set the elasticity of intertemporal substitution (EIS) to generate the same peak response of output to an identified monetary shock as in the HANK model. In practice, the EIS is little changed.

agree on the transmission from interest rate changes to output and inflation, the interest rate movements that achieve the policymaker's desired output and inflation movements cannot be too dissimilar.[14]

5 Optimal Ramsey policy

We now turn to the optimal policy problem of a Ramsey planner. Unlike our *ad hoc* dual-mandate loss function of Section 4, this planner's objective is directly affected by household inequality, reflecting a desire to dampen the distributional consequences of aggregate shocks. For most of this section we will focus on optimal monetary policy. We will, however, also discuss *joint* optimal monetary and stimulus check policy, emphasizing the complementarity between these two standard stabilization tools.

The remainder of this section proceeds in three steps. First, in Section 5.1, we show how to express the optimal policy problem in our general linear-quadratic form. Second, in Section 5.2 we present general analytical results and discuss some instructive analytical special cases. Finally, in Sections 5.3 and 5.4, we turn to quantitative analysis, first connecting our

[14]As discussed above, the two interest rate paths would agree exactly if the two models were to agree on all of $\Theta_{y,i}$ and $\Theta_{\pi,i}$. Our calibration instead only ensures that peak output and inflation responses in response to a particular, transitory interest rate movement align. This limits the extent of—but does not fully eliminate—disagreement in nominal rate paths.

optimal rule expressions to empirical evidence, and then applying our results to particular business-cycle shocks.

5.1 The optimal policy problem

We consider a conventional Ramsey planner that aggregates utilities of the households populating the economy. This section presents a linear-quadratic version of this policy problem. As is typical in the optimal policy literature (e.g. see Giannoni & Woodford, 2002), doing so requires a particular notion of efficiency of the steady state—here efficiency in the sense that the planner finds the steady-state cross-sectional distribution of consumption desirable. Our optimal policy analysis will thus reflect an insurance motive against fluctuations in these consumption shares.

Loss function. To state the loss function, it will prove convenient to describe an individual's outcomes in terms of their idiosyncratic history of 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.[15] Letting $\Gamma(\zeta)$ denote the (stationary) distribution of such histories (with ζ a generic realization of a history), we can write the social welfare function as

$$\mathcal{V}^{HA} = \sum_{t=0}^{\infty} \beta^t \int \varphi(\zeta) \left[\frac{(\omega_t(\zeta)c_t)^{1-\gamma} - 1}{1-\gamma} - \nu(\ell_t) \right] d\Gamma(\zeta), \tag{28}$$

where $\varphi(\zeta)$ is a Pareto weight on the utility of households with history ζ.

In keeping with optimal (monetary) policy analysis in standard representative-agent environments (Woodford, 2003), our objective is to evaluate the social welfare function (28) to second order. To this end, a first-order approximation to aggregate equilibrium dynamics suffices only if the expansion point (i.e., the deterministic steady state) is efficient. Without household heterogeneity, a simple production subsidy is sufficient to ensure this. With household heterogeneity, however, we now additionally require the consumption *shares* of all households to be optimal. In principle there are two ways of ensuring this: either the steady-state fiscal tax-and-transfer system achieves the optimal level of insurance given the planner weights $\varphi(\bullet)$, or the planner weights are set residually so that the implied steady-

[15]Note that this is without loss of generality, as individuals in our model are *ex ante* identical, so their outcomes only differ due to different histories of shocks.

state distribution of consumption given a tax-and-transfer system is optimal. We adopt the second approach, following much of the inverse optimal taxation literature (e.g. Heathcote & Tsujiyama, 2021). Our preference for this approach reflects the overarching focus of this paper: we ask how cyclical policy tools should be manipulated to respond to cyclical changes in inequality, leaving the long-run steady state outside of the purview of our analysis.

Appendix D.3 presents our assumptions on the production subsidy and policymaker preference weights that ensure efficiency of the deterministic steady state. Given those assumptions, a second-order approximation of (28) around the efficient steady state then gives the following characterization of the policymaker loss function.

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

$$\mathcal{L}^{HA} \equiv \sum_{t=0}^{\infty} \beta^t \left[\widehat{\pi}_t^2 + \frac{\kappa}{\bar{\eta}} \widehat{y}_t^2 + \frac{\kappa\gamma}{(\gamma+\phi)\bar{\eta}} \int \frac{\widehat{\omega}_t(\zeta)^2}{\bar{\omega}(\zeta)} d\Gamma(\zeta) \right], \tag{29}$$

where $\widehat{\omega}_t(\zeta) = \omega_t(\zeta) - \bar{\omega}(\zeta)$ and $\bar{\omega}(\zeta)$ is the steady-state consumption share of an individual with history ζ.

Note that, in the representative-agent analogue of our economy (as discussed in Section 3.5), the loss function would feature the same first two terms, as already well-known from prior work (Woodford, 2003). Our analysis reveals that household heterogeneity adds a third, inequality-related term, with the planner wishing to stabilize the consumption *shares* of everyone in the economy.

How does the inequality term in (29) fit into the linear quadratic framework in Section 2? Moving to a sequence-space formulation, we can write the loss as

$$\mathcal{L}^{HA} = \lambda_\pi \widehat{\boldsymbol{\pi}}' W \widehat{\boldsymbol{\pi}} + \lambda_y \widehat{\boldsymbol{y}}' W \widehat{\boldsymbol{y}} + \int \lambda_{\omega(\zeta)} \widehat{\boldsymbol{\omega}}(\zeta)' W \widehat{\boldsymbol{\omega}}(\zeta) d\Gamma(\zeta), \tag{30}$$

where $\lambda_\pi = 1$, $\lambda_y = \frac{\kappa}{\bar{\eta}}$ and $\lambda_{\omega(\zeta)} \equiv \frac{\kappa\gamma}{(\gamma+\phi)\bar{\eta}\bar{\omega}(\zeta)}$. The consumption share for each idiosyncratic history thus emerges as a separate target variable for the policymaker. We will discuss our approach to computation of (30) later in this section.[16]

Constraints. 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

[16]Note that, technically, (30) does not immediately fit into our framework in Section 2.1 since the objective here features an integral (rather than a simple sum). Our approach to computation will consider an equivalent formulation of the problem with finitely many policy targets.

shocks and policy choices. As already discussed in Section 4.1, the evolution of output and inflation are governed by the Phillips curve (17) and the IS curve (18). It then remains to describe the evolution of the inequality term in (29). In Appendix C.2, we establish that, to first order, we can write the consumption share for a household with specific history ζ as

$$\widehat{\boldsymbol{\omega}}(\zeta) = \Omega_{\omega(\zeta),y}\widehat{\boldsymbol{y}} + \Omega_{\omega(\zeta),\pi}\widehat{\boldsymbol{\pi}} + \Omega_{\omega(\zeta),i}\widehat{\boldsymbol{i}} + \Omega_{\omega(\zeta),x}\widehat{\boldsymbol{\tau}}_x + \Omega_{\omega(\zeta),m}\widehat{\boldsymbol{m}}, \qquad \forall\zeta \tag{31}$$

where the maps $\Omega_\bullet$ give the derivatives of consumption shares with respect to aggregate variables. The intuition for (31) 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. By the proof of Lemma 1, we can obtain these inputs as a function of exogenous shocks, policies, and aggregate output and inflation paths.

Following McKay & Wolf (2021), we can alternatively re-write this constraint in impulse response space, solving out the dependence of consumption shares on income and inflation:

$$\widehat{\boldsymbol{\omega}}(\zeta) = \Theta_{\omega(\zeta),i}\widehat{\boldsymbol{i}} + \Theta_{\omega(\zeta),x}\widehat{\boldsymbol{\tau}}_x + \Theta_{\omega(\zeta),\eta}\widehat{\boldsymbol{\eta}} + \Theta_{\omega(\zeta),m}\widehat{\boldsymbol{m}}. \qquad \forall\zeta \tag{32}$$

Summary & computational details. The Ramsey planner chooses paths of the two available policy instruments—nominal interest rates $\boldsymbol{i}$ and the exogenous component of transfers $\boldsymbol{\tau}_x$—to minimize the derived loss function $\mathcal{L}^{HA}$ subject to the same two constraints as before, (17) and (18), as well as the evolution of the inequality term, (31).

To computationally evaluate the more complicated loss function (30) and the associated constraints, we show in Appendix C.2 that the Ramsey loss can be re-written as

$$\mathcal{L}^{HA} = \widehat{\boldsymbol{\pi}}'W\widehat{\boldsymbol{\pi}} + \frac{\kappa}{\bar{\eta}}\widehat{\boldsymbol{y}}'W\widehat{\boldsymbol{y}} + \frac{\kappa\gamma}{(\gamma+\phi)\bar{\eta}}\widehat{\boldsymbol{x}}'Q\widehat{\boldsymbol{x}} \tag{33}$$

where $\boldsymbol{x} = (\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\tau}_x, \boldsymbol{\tau}_e, \boldsymbol{m})$ and Q is a linear map, defined in Appendix C.2.[17] The alternative formulation in (33) reflects the simple intuition that it is always possible to re-write the loss coming from cross-sectional inequality as a function of the (small number of) inputs to the household consumption-savings problem—income, interest rates, taxes, and shocks. With the re-written loss function (33), the relevant constraints are then simply the equilibrium dynamics of $\boldsymbol{x} = (\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\tau}_x, \boldsymbol{\tau}_e, \boldsymbol{m})$. Computation is thus again straightforward: Appendix C.2

[17]The linear map Q in (33) is in general not diagonal. Thus Appendix C.2 also extends the results from Section 2 to the case with such interaction terms.

discusses how to recover the map Q, while the coefficient matrices for all constraints are again immediate from standard sequence-space solution output, exactly as in the dual-mandate problem considered before. To summarize, relative to solving a HANK model *given* a policy rule, the only additional computational work needed to solve a Ramsey *optimal* policy problem is the one-time computation of the auxiliary matrix Q.

5.2 Optimal policy rules

Having expressed the optimal Ramsey problem in linear-quadratic form, we can now leverage the results of Section 2 to provide a general characterization of optimal Ramsey policy rules.

The optimal Ramsey monetary policy rule is given as

$$\Theta'_{\pi,i} W \widehat{\boldsymbol{\pi}} + \frac{\kappa}{\bar{\eta}} \Theta'_{y,i} W \widehat{\boldsymbol{y}} + \int \lambda_{\omega(\zeta)} \Theta'_{\omega(\zeta),i} W \widehat{\boldsymbol{\omega}}(\zeta) d\Gamma(\zeta) = \mathbf{0}, \tag{34}$$

where as before the matrices $\Theta_{\bullet,i}$ collect the dynamic causal effects of interest rate movements on the various policymaker targets. In particular, we see that the first two terms in (34) are identical to the optimal dual-mandate rule (25), just now with weights derived from 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 rule

$$\Theta'_{\pi,\tau_x} W \widehat{\boldsymbol{\pi}} + \frac{\kappa}{\bar{\eta}} \Theta'_{y,\tau_x} W \widehat{\boldsymbol{y}} + \int \lambda_{\omega(\zeta)} \Theta'_{\omega(\zeta),\tau_x} W \widehat{\boldsymbol{\omega}}(\zeta) d\Gamma(\zeta) = \mathbf{0}, \tag{35}$$

As discussed in Wolf (2021) and in Section 4.2, interest rate and stimulus check policy in our environment implement the same output-inflation allocations, so the first two terms in (34) and (35) 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, thereby opening the door for interest rate and stimulus check policies to be useful complementary tools for aggregate stabilization. We will return to this observation in Section 5.4. Overall, (34) and (35) fully characterize joint optimal monetary-fiscal policy.

The key takeaway from the characterizations in (34) and (35) is that optimal policy rules in the general Ramsey problem deviate from the dual-mandate rules discussed in Section 4 if and only if the policy instruments affect cross-sectional consumption inequality. On the one hand, if monetary policy is distributionally neutral—as is for example the case in the environment of Werning (2015) (see Appendix A.2 for details)—then household inequality does not affect the optimal rule: $\Theta_{\omega(\zeta),i} = \mathbf{0}$ implies that the optimal monetary rule continues to take the dual-mandate form[18]

$$\widehat{\pi}_t + \frac{1}{\bar{\eta}} \left(\widehat{y}_t - \widehat{y}_{t-1}\right) = 0, \quad \forall t = 0, 1, \dots \tag{36}$$

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 (29) may materially change policy conduct, with insurance concerns swamping the usual price and output stabilization motives. For example, in response to a cost-push shock that redistributes income from workers to capitalists, the central bank may be reluctant to aggressively hike rates, to avoid further worsening consumption inequality. Ultimately, which case is the relevant one is an inherently empirical question.

Our strategy. The optimal policy analysis in the remainder of this paper will proceed by confronting the general expressions in (34) and (35) with empirical evidence. Given the paper's emphasis on monetary policy, we pay particular attention to evidence on the response of consumption inequality to changes in nominal rates—that is, entries of the causal effect matrices $\Theta_{\omega(\zeta),i}$. Section 5.3 begins by reviewing that empirical evidence, before then using it to calibrate our structural model. Applications follow in Section 5.4.

5.3 Empirical evidence, calibration, and model validation

We have seen that the extent to which the insurance motive embedded in (29) affects optimal policy design depends crucially on the distributional effects of changes in policy instruments. By the results in McKay & Wolf (2021), we can assess these distributional effects by studying the implications of identified policy shocks.

In this section we do so by proceeding in three steps. First, we review the empirical evidence on the distributional effects of monetary policy shocks. Second, we calibrate our structural HANK model to be consistent with that evidence. And third, we validate our model against untargeted moments.

[18] Bilbiie (2021) 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.

The distributional effects of monetary policy. We are ultimately interested in the effects of policy changes on consumption inequality. Empirical work, however, has so far provided the most definitive answers not on consumption itself, but on inputs to the consumption-savings decision—that is, on various *channels* of monetary policy propagation. We begin with a brief review of the evidence on the most important channels.

a) *Labor income.* There is extensive evidence that those with less education, low past earnings, and racial minorities tend to be more exposed to cyclical fluctuations in labor market conditions (e.g. Okun, 1973; Hoynes, 2000; Guvenen et al., 2014; Patterson, 2022). Turning to monetary policy more specifically, Guvenen et al. (2014) study earnings dynamics in the 1979–1983 recession—a contraction that was arguably caused by a large monetary intervention. They find that individuals with low previous earnings also suffered the largest earnings losses. Several recent studies that combine identified monetary policy shocks with European administrative data arrive at very similar conclusions (Andersen et al., 2021; Amberg et al., 2021; Holm et al., 2021).

b) *Asset prices.* Expansionary monetary policy raises the value of long-duration assets and leads to considerable capital gains for the owners of those assets. For example, Bauer & Swanson (2022) estimate that a 100 basis point reduction in short-term interest rates increases the S&P 500 stock market index by about five percentage points. As the distribution of wealth is concentrated, these capital gains are very unequally distributed across the population. Andersen et al. (2021) leverage administrative household-level data from Denmark to quantify this distributional gradient. They find that, in response to a one percentage point decrease in the policy rate, asset values increase by around 70 per cent of total annual disposable household income at the top of the income distribution, and around 30 per cent at the median.[19]

c) *Nominal wealth redistribution.* Surprise inflation will lower the real value of nominal positions. Doepke & Schneider (2006) report that wealthy households have about 10-15% of their net worth in nominal positions, while young middle-class households have considerable nominal liabilities in the form of mortgages. The magnitude of the redistribution

[19] Of course, while long-duration real assets increase in value following an expansionary monetary shock, so too do the values of long-duration liabilities. In particular, the future consumption plans of households become more costly as real interest rates fall. Theoretically, one would want to measure the net exposure to interest rates that accounts for the difference between asset income and planned consumption (see e.g. Auclert, 2019).

that occurs through this channel of course depends on the strength of the inflation response to monetary policy. Most empirical studies point to a relatively flat Phillips curve (see e.g. Mavroeidis et al., 2014), thus implying relatively small inflation responses to changes in policy, and so only moderate nominal wealth redistribution.

d) *Mortgage payments.* Expansionary policy lowers mortgage interest rates, thus leading to lower debt service payments for households that buy a new home or choose to refinance an existing mortgage. This channel most strongly benefits middle-class households that own a home with a large mortgage relative to their net worth; renters and homeowners who own their home outright, on the other hand, are not directly affected (e.g. see Cloyne et al., 2020; Wong, 2021).

Several studies have tried to go beyond these inputs to the consumption-savings decision and examine the response of household consumption directly. Our preferred estimates come from Holm et al. (2021)—to the best of our knowledge the only study in which households are tracked over time and household fixed effects are included. One of their main findings is that expansionary policy has U-shaped effects on household consumption across the distribution of liquid asset holdings.[20] Other work however arrives at somewhat different conclusions: using data from the Consumer Expenditure Survey, Coibion et al. (2017) find that expansionary policy leads consumption inequality to fall, while Chang & Schorfheide (2022) conclude that inequality increases, reflecting a large increase in consumption in the right tail.

Overall, our reading of the literature is that conclusions on *channels* are more definitive than conclusions on the implied consumption movements. Our strategy is thus to match the empirical evidence on channels of redistribution in our model, and then infer the implications for consumption, leveraging the fact that the canonical incomplete-markets model is a successful model of the consumption-savings problem (e.g. Kaplan et al., 2018).[21]

Model calibration. We calibrate our model to capture the just-summarized channels of monetary policy propagation to consumption inequality. Since our model does not feature mortgages, we only match channels a)–c), and so our model may understate the benefits of

[20]Like many empirical studies of monetary policy shocks, Holm et al. find that monetary policy affects the labor market with a considerable delay. Therefore the induced changes in non-financial income also occur with a delay. As it takes time for these indirect effects of monetary policy to materialize we focus on the consumption response at horizons of two to five years.

[21]A similar strategy of using a model of consumption and savings decisions to infer the implications of monetary shocks for the consumption distribution has been employed by Ampudia et al. (2018). They too find U-shaped effects on consumption.

expansionary policy to middle class households. As we find that these households benefit the least, actual heterogeneity in the distributional effects of monetary policy may be even less pronounced than our model implies, thus reinforcing our conclusions in Section 5.4.

- *Income process.* Since many of the distributional effects of monetary policy involve changes in asset values, it is paramount that the model generates a concentrated distribution of wealth, as in the data. Since standard incomplete-markets models of the consumption-savings decision struggle to generate sufficiently concentrated wealth holdings in the top tail of the distribution, we follow Castaneda et al. (2003), Boar & Midrigan (2020), and Greenwald et al. (2021) in specifying an income process with superstar earners, allowing us to generate a realistic concentration of wealth.

  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

$$\log e_{it} = \log(\zeta_{it})(1 + m_t + \alpha \log y_t) + \log \bar{e}_t, \tag{37}$$

  where α controls the sensitivity of income dispersion to the cycle. A negative α implies that low-ζ households are more exposed to the cycle. $\bar{e}_t$ is a normalization constant such that $\int e_{it} di = 1$ at all dates. For regular workers, the income state $\log \zeta_{it}$ follows an AR(1) process. To estimate this process we follow Guvenen et al. (2022). Those authors propose a parametric income process that allows for differential exposure to aggregate conditions, and then estimate it using moments taken from the Social Security Administration data. Adopting this strategy allows us to pin down α in (37). Alternatively, households can be high earners, with households entering and exiting the high-earnings state at constant rates. We allow for two levels of high earnings, and then calibrate the size of these groups, the level of their earnings, and the persistence of these states to match data on the income and wealth distributions. Specifically, we target the shares of total wealth held by the wealthiest 1%, 5%, 10%, 25% and 50% of the wealth distribution. We also target the same moments of the income distribution. Table 1 shows the targets and fitted values. Further details on the calibration of the income process and its parameters are presented in Appendix B.1.

- *Asset structure.* The steady state real interest rate $\bar{r}$ is set to 2 per cent annually. The duration of the long-run bond is set to 15 years, following Greenwald et al. (2021). Next, we set $\chi_\pi = 0.9$ in reference to the estimates for wealthy households in Doepke &

	Wealth		Income	
	Data	Model	Data	Model
Top 1%	37	36	17	20
Top 5%	65	66	32	36
Top 10%	76	78	43	52
Top 25%	91	95	64	65
Top 50%	99	100	84	80

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

Schneider (2006), while χ_y is calibrated internally to match the sensitivity of aggregate household net worth to an identified monetary shock. Details on the internal calibration are provided in Appendix B.1. The aggregate supply of assets is set to match the average ratio of household net worth to GDP (with household net worth computed from the U.S. Financial Accounts), and the borrowing limit $\underline{a}$ is set to zero, as commonly done in the literature (e.g. McKay et al., 2016).

- *Fiscal system.* The U.S. fiscal tax-and-transfer system is reasonably well-approximated by a common 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.[22] The speed of fiscal adjustment is controlled by σ, which we set equal to δ. Bonds in excess of the steady state level are thus paid off as they mature.[23]

- *Other parameters.* We set the Frisch elasticity to one, and the elasticity of substitution between labor varieties to six, based on Basu & Fernald (1997). Lastly, we set the coefficient 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. Again, details for this internal calibration are provided in Appendix B.1.

We summarize all parameter values in Table 2.

[22] 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 the steady state transfer.

[23] We have found that the results are robust to the choice of σ as long as it takes a low value. Empirically, changes in public debt are highly persistent implying a low value of σ.

Parameter	Description	Value	Calibration target
Households			
$\{\zeta_{it}\}$	Income risk process	–	See text
α	Income exposure	-1.61	Het. earnings cyclicality
γ	Relative risk aversion	1.5	Monetary shock effects
ϕ	Frisch elasticity	1	Standard
β	Discount factor	0.985	Asset market clearing
Firms			
κ	Phillips curve slope	0.027	Monetary shock effects
$\bar{\eta}$	Labor Substitutability	6	Basu & Fernald (1997)
Asset Structure			
$\bar{r}$	Steady-state return	0.5%	2% annual real return
$\underline{a}$	Borrowing limit	0	Standard
$\frac{\bar{a}}{\bar{y}}$	Total asset supply	16.24	Household net worth/GDP
δ	Asset duration	0.048	Greenwald et al. (2021)
χ_π	Inflation indexation	0.9	Doepke & Schneider (2006)
χ_y	Output sensitivity	7.0	Monetary shock effects
Government			
τ_y	Labor tax	0.248	Steady-state budget balance
$\frac{\bar{\tau}_x}{\bar{y}}$	Transfer share	0.17	After-tax vs. before-tax income
σ	Tax-Debt Responsiveness	$= \delta$	See text

Table 2: Calibration of our quantitative HANK model. The model period is one quarter.

Consumption responses & model validation. Having calibrated our model to be consistent with empirical evidence on the transmission channels of monetary policy to households, we now use the model to map these income responses into monetary policy's effects on consumption inequality.

Our main finding is that the consumption of low-wealth and high-wealth households is more sensitive to monetary policy than the consumption of households with moderate assets. Figure 2 provides an illustration. The figure shows the initial change in consumption across the wealth distribution following an expansionary monetary shock. On the one end of the spectrum, low-wealth households tend to have low incomes, their earnings are more exposed to aggregate income, and they are often borrowing-constrained and so have high MPCs. As a result, changes in earnings pass through strongly into consumption. At the other end, high-wealth households benefit from increases in asset values. While these capital gains are very large, the pass-through to consumption is relatively weak. Households with a moderate amount of wealth are less affected, thus giving the U-shape displayed in Figure 2.[24]

Finally, Figure 2 also provides an important validation check of our model. In the figure, the two black lines show the consumption response to a monetary shock as estimated by Holm et al. (2021). Since empirical estimates suggest that monetary policy transmits with a lag, we follow Holm et al. and focus on the consumption response at horizons of 2 to 5 years. Holm et al. classify households by their liquid asset holdings; assuming that liquid assets are monotonic in wealth, we can compare their results to our model's implications. As Figure 2 shows, model and data both qualitatively and quantitatively agree on the cross-sectional distributional effects of a monetary easing on household consumption.

5.4 Applications

With a model that matches empirical evidence on the distributional effects of monetary policy in hand, we now use this laboratory to explore optimal Ramsey policy in two empirically relevant scenarios. As our first application, we consider 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. Second, we return to the inflationary cost-push

[24]Note that the earnings of the lowest-wealth group increase by 0.5% of their steady-state income, while the capital gains of the highest-wealth group amount to more than 20% of steady-state income. The model is thus consistent with the empirical results of Andersen et al. (2021) and Bartscher et al. (2021) who emphasize that capital gains effects are large relative to labor income effects. While the implications for short-run consumption are more equal, high-wealth consumption does remain elevated for longer. These dynamic effects are not visible in Figure 2, but they are incorporated into our analysis of optimal policy.

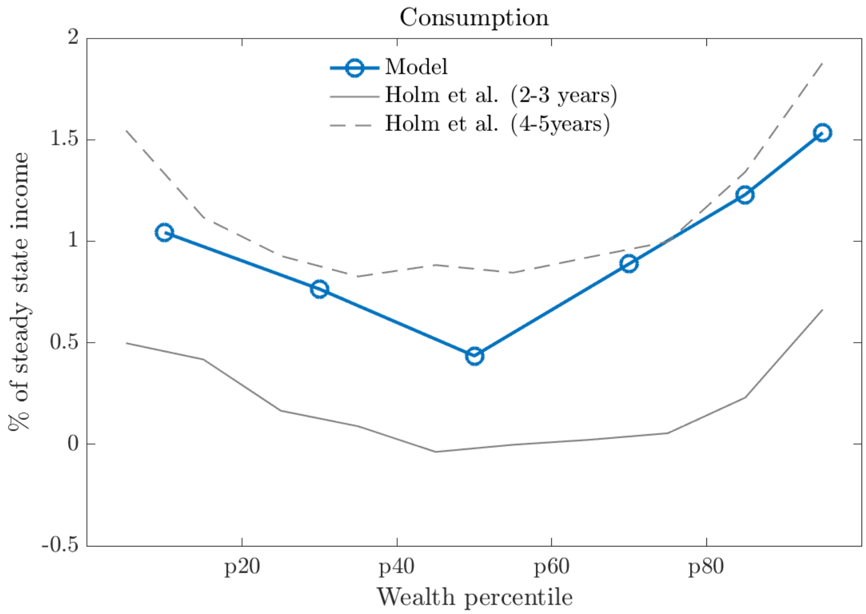

Figure 2: Initial response of consumption to an expansionary monetary policy shock across the distribution of wealth. The empirical estimates are from Holm et al. (2021), who rank households according to liquid assets, which we assume are monotonic in wealth in constructing this figure. We simulate the shock we estimate empirically in Appendix B.1 scaled to match the magnitude of the consumption responses in Holm et al.. Holm et al. find that the indirect effects of policy build through time whereas in our model they occur on impact.

shock from Section 4.3—a shock that has received much attention in the HANK optimal policy literature (e.g. Bhandari et al., 2021; Acharya et al., 2020).

Distributional shock. Our first shock is an innovation to m_t, the exogenous driver of income dispersion in (37). This shock redistributes income from low-income households to high-income households, thus depressing aggregate demand: precautionary savings increase due to the increase in risk, and spending falls as income is redistributed towards lower-MPC households. We study three optimal policy responses to this shock: monetary policy for a dual mandate policymaker; monetary policy for a Ramsey planner; and joint monetary-fiscal policy for a Ramsey planner. Results are reported in Figure 3.

As a benchmark we begin with the optimal dual mandate policy response, displayed as the grey lines in Figure 3. As aggregate demand falls, the dual-mandate central banker cuts nominal interest rates to perfectly stabilize output and inflation ("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 the monetary

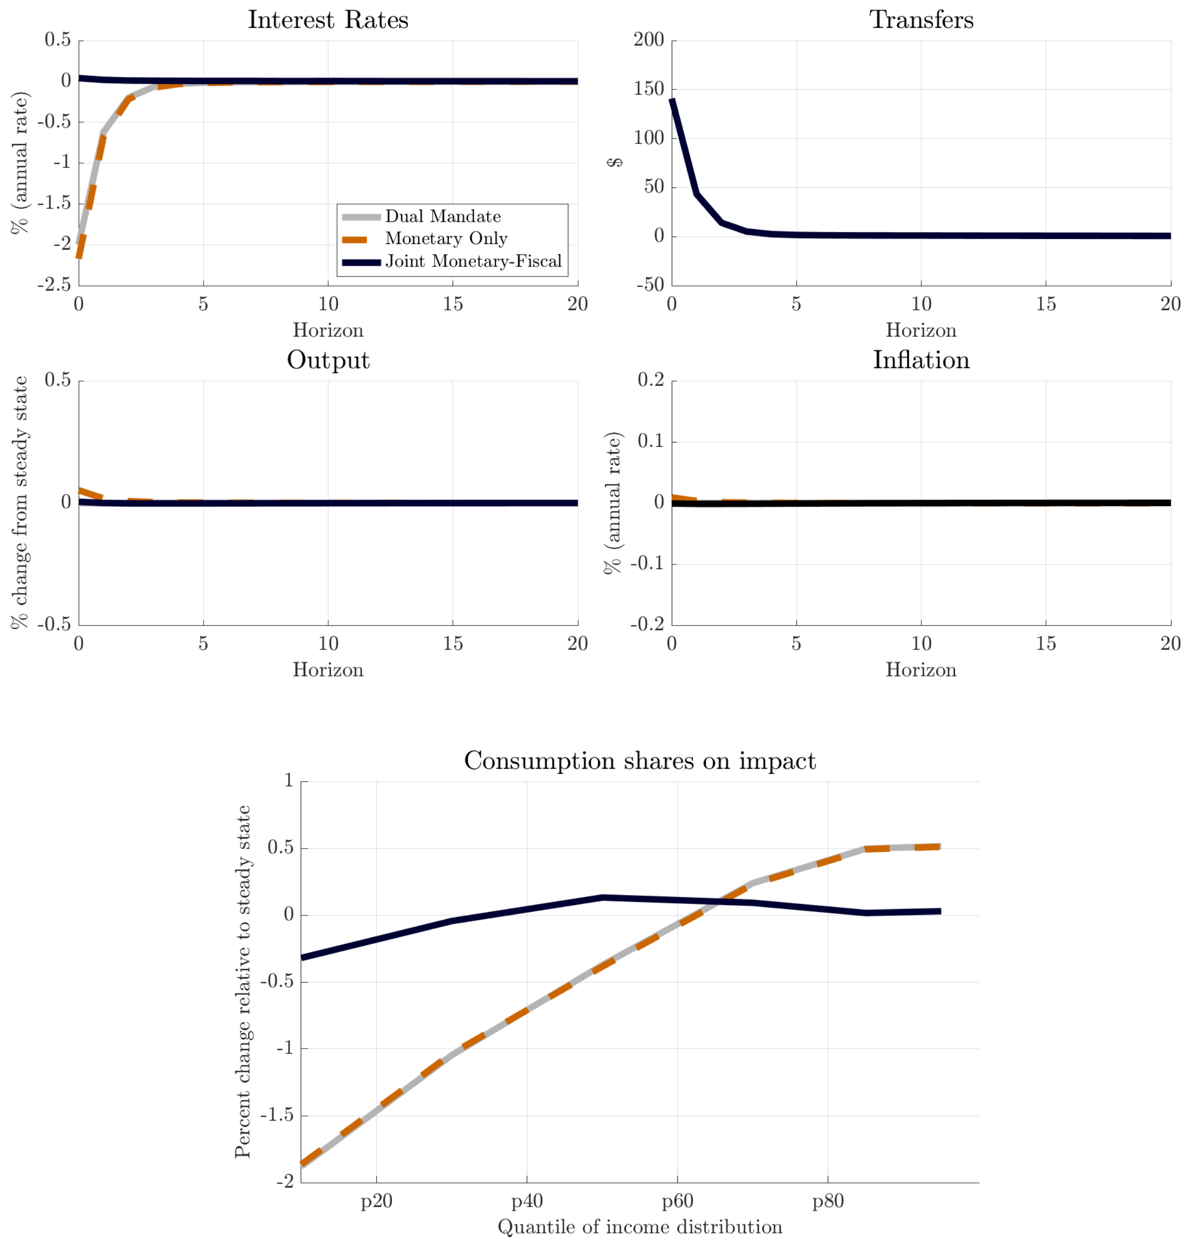

Figure 3: Optimal policy response to an income distribution shock m_t. The figure shows the results from three policy rules: optimal monetary policy for a dual-mandate policymaker (Dual Mandate), optimal monetary policy for a Ramsey planner that is constrained not to use lump-sum transfers (Monetary Only), and optimal policy for a Ramsey planner that can use both tools (Joint Monetary-Fiscal). Transfers are expressed in units of dollars, using the conversion that steady state per capita GDP is $60,000.

easing only moderately offsets these effects, somewhat stabilizing consumption at the bottom but increasing consumption at the top even further.

We then consider the optimal Ramsey monetary policy that also considers distributional objectives, depicted as the orange dashed lines in Figure 3. Our headline finding here is that the optimal policy response is very similar to the dual-mandate policy, with interest rates cut only slightly more (17 basis points). As a result, output and inflation continue to be stabilized almost perfectly, and the consumption distribution hardly differs from the dual-mandate outcome. The intuition is as follows. The original shock m_t has a strong distributional tilt, with low-income households losing the most. The Ramsey planner would like to lean against this redistribution and stabilize consumption shares. Monetary policy, however, is a rather blunt tool: if the planner were to cut rates by enough to stabilize consumption at the bottom, then consumption of the rich and so aggregate output and inflation would overshoot significantly—an immediate implication of the distributional effects we matched in Figure 2. We provide a visual illustration of this point in Figure B.4: approximately doubling the size of the nominal interest rate cut is enough to nearly stabilize the consumption of households with low income and low wealth, but it comes at the cost of over-stimulating the consumption of high-wealth households (and thus aggregate output and inflation).

Finally we turn to the optimal joint monetary-fiscal Ramsey policy, with a planner using both interest rate policy and stimulus checks to pursue her aggregate and distributional objectives. Results are displayed as the dark blue lines in Figure 3. Compared to monetary policy, fiscal stimulus payments are much more progressive, with consumption of low-income and low-wealth households responding significantly more than that of high-income, high-wealth households (see Wolf (2021) or Figure B.3). Intuitively, this reflects both differences in MPCs as well as differences in income levels, with any given dollar amount of 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 differently affect low-income and high-income households. The results depicted in Figure 3 are consistent with this intuition: in response to the distributional shock m_t, the Ramsey planner can use stimulus 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 for the monetary authority to intervene, with nominal interest rates now responding very little.

Overall, our analysis of optimal Ramsey policy responses to the distributional shock suggests two broad lessons for optimal cyclical policy design. First, for monetary policy viewed in isolation, an important question is how well-adapted its distributional effects are likely to be to the distributional incidence of the underlying business-cycle shock. In many cases, stabilizing consumption at the bottom of the distribution may require significant departures from dual-mandate objectives; if the Ramsey planner is unwilling to accept these departures (as was the case in our example), then optimal policy will stay reasonably close to the dual-mandate benchmark. Second, fiscal stimulus payments and interest rate policy can be highly complementary policy tools. They are likely to have very different cross-sectional incidence profiles (see Figure 2 and Figure B.3), and so will be well-suited to offset cyclical shocks with similarly distinct distributional incidence profiles.

Cost-push shock. Our second shock is a cost-push shock η_t, as already considered in Section 4.3. In our model, this cost-push shock introduces a wedge in the Phillips curve, thus giving an output-inflation trade-off. The shock does not, however, by itself have strong distributional implications. For our analysis we restrict attention to monetary policy alone, comparing only the optimal dual-mandate and Ramsey monetary policy responses.[25] Results are displayed in Figure 4.

Overall we find that the Ramsey policy responds less aggressively than the dual-mandate policy. The intuition is as follows. The increase in nominal interest rates that is optimal under a dual mandate objective by itself leads to an inverse-U-shaped effect on consumption. To mitigate these distributional effects, the optimal Ramsey policy response is somewhat attenuated relative to the dual mandate benchmark. We conclude that insurance concerns in this experiment do affect optimal monetary policy design, with the initial decline in output 44 per cent smaller than under the dual-mandate rule.

However, we emphasize that the broad thrust of the policy response—raising rates to lean against inflation—still follows the prescription of standard representative-agent models. These quantitative findings contrast with recent arguments made for example in Bhandari et al. (2021), who conclude that distributional considerations may lead to substantial changes in optimal policy conduct—changing the signs of impulse response functions and the order of magnitude of inflation volatility. The reason for this disagreement is that, in the framework of Bhandari et al., expansionary monetary policy strongly redistributes from wealthy households to poor households, thus making it a very useful tool for achieving distributional

[25] The results for joint optimal monetary-fiscal policy are anyway not particularly different. The intuition is that the original cost-push shock does not have a particularly pronounced cross-sectional incidence profile. As discussed below, the optimal monetary response creates distributional effects, but stimulus checks are of course ill-suited to offset these effects.

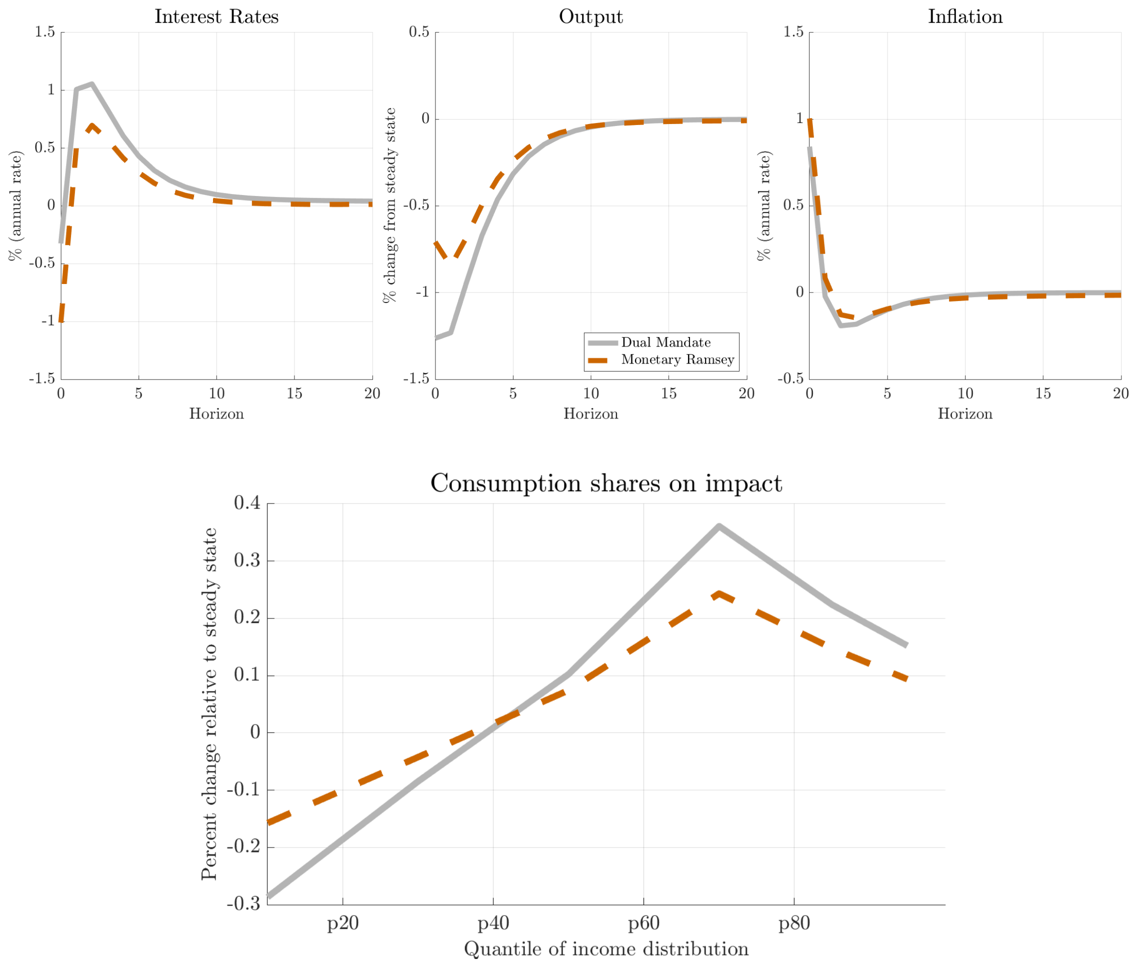

Figure 4: Optimal policy response to an inflationary cost-push shock. The figure shows optimal monetary policy for a planner whose objective is to stabilize output and inflation (Dual Mandate) and for a Ramsey planner (Monetary Ramsey).

goals.[26] The recent empirical evidence on the distributional effects of monetary policy reviewed in Section 5.3 on the other hand suggests that these distributional effects are rather muted, leading us to find a much smaller departure from dual-mandate policy.

6 Conclusion

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

[26]This redistribution occurs primarily though raising the labor share and reducing monopoly profits in their sticky price model.

First, for central banks that target standard macroeconomic aggregates (e.g., a classical "dual mandate"), household inequality is likely to only have moderate effects. Analytically, we have given conditions under which the optimal forecast target criterion of a dual mandate central banker is unaffected by household inequality. Under those conditions, optimal output and inflation outcomes will be completely independent of household inequality. 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 furthermore yield similar predictions for the paths of nominal rates necessary to implement a given output and inflation target—irrespective of whether the model features household inequality or not.

Second, the extent to which distributional objectives shape optimal Ramsey monetary policy depends crucially on the causal effects of interest rate changes on household inequality. According to our reading of the empirical evidence, monetary policy indeed has meaningful distributional effects, but these effects are not straightforward—in response to a monetary easing, both poor as well as rich households are likely to gain the most. 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, fiscal stimulus checks—an alternative tool of stabilization policy, used increasingly frequently in recent decades—promises to be more useful for distributional purposes. Such stimulus checks achieve aggregate stabilization *through* insurance at the bottom, thus making them well suited to address cyclical fluctuations that mostly affect poor households.

References

Acharya, S., Challe, E., & Dogra, K. (2020). Optimal monetary policy according to hank. CEPR Discussion Paper No. DP14429.

Amberg, N., Jansson, T., Klein, M., & Rogantini Picco, A. (2021). Five facts about the distributional income effects of monetary policy. CESifo Working Paper.

Ampudia, M., Cooper, R., Le Blanc, J., & Zhu, G. (2018). Mpc heterogeneity in europe: Sources and policy implications. NBER Working Paper 25082.

Andersen, A. L., Johannesen, N., Jørgensen, M., & Peydró, J.-L. (2021). Monetary policy and inequality. Working Paper.

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

Auclert, A., Bardóczy, B., & Rognlie, M. (2020). Mpcs, mpes and multipliers: A trilemma for new keynesian models. Technical report, National Bureau of Economic Research.

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. (2018). The intertemporal keynesian cross. Technical report, National Bureau of Economic Research.

Bartscher, A. K., Kuhn, M., Schularick, M., & Wachtel, P. (2021). Monetary policy and racial inequality. CEPR Discussion Paper No. DP15734.

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

Bauer, M. D. & Swanson, E. T. (2022). A reassessment of monetary policy surprises and high-frequency identification. National Bureau of Economic Research Working Paper.

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

Bernanke, B. (2015). Remarks at Brookings Institution: Objections to Federal Reserve Accountability Bill.

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. (2021). Monetary policy and heterogeneity: An analytical framework. Working Paper.

Bilbiie, F. O., Monacelli, T., & Perotti, R. (2021). Stabilization vs. redistribution: the optimal monetary-fiscal mix. CEPR Discussion Paper No. DP15199.

Boar, C. & Midrigan, V. (2020). Efficient redistribution. National Bureau of Economic Research Working Paper.

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., Harbo Hansen, N.-J., Krusell, P., & Öberg, E. (2020). The new keynesian transmission mechanism: A heterogeneous-agent perspective. *The Review of Economic Studies*, *87*(1), 77–101.

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.

Cloyne, J., Ferreira, C., & Surico, P. (2020). Monetary policy when households have debt: new evidence on the transmission mechanism. *The Review of Economic Studies*, *87*(1), 102–129.

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.

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

Doepke, M. & Schneider, M. (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy*, *114*(6), 1069–1097.

Farhi, E. & Werning, I. (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, *109*(11), 3887–3928.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., & Schorfheide, F. (2016). Solution and estimation methods for DSGE models. In *Handbook of Macroeconomics*, volume 2 (pp. 527–724). Elsevier.

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.

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.

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

Hebden, J. & Winker, F. (2021). Impulse-Based Computation of Policy Counterfactuals. 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. (2021). Should monetary policy care about redistribution? optimal fiscal and monetary policy with heterogeneous agents. Technical report, Working Paper, SciencesPo.

Mavroeidis, S., Plagborg-Møller, M., & Stock, J. H. (2014). Empirical evidence on inflation expectations in the new keynesian phillips curve. *Journal of Economic Literature*, *52*(1), 124–88.

McKay, A., Nakamura, E., & Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, *106*(10), 3133–58.

McKay, A. & Wolf, C. K. (2021). What can time-series regressions tell us about policy counterfactuals? MIT Working Paper.

Nuno, G. & Thomas, C. (2021). Optimal redistributive inflation. Working Paper.

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

Patterson, C. (2022). The matching multiplier and the amplification of recessions. Working Paper.

Plagborg-Møller, M. & Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, *89*(2), 955–980.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. *Handbook of macroeconomics*, *2*, 71–162.

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. (2021). Interest rate cuts vs. stimulus payments: An equivalence result. National Bureau of Economic Research Working Paper 29193.

Wong, A. (2021). Refinancing and the transmission of monetary policy to consumption. Working paper.

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

A Supplementary model details

This appendix provides further details for our structural model. Appendix A.1 begins by further discussing the union problem and deriving the log-linearized Phillips curve (11), and Appendix A.2 discusses the model's special case in which monetary policy is distributionally neutral, following Werning (2015).

A.1 Technology, union problem & Phillips Curve

We here provide further details for the production side of our economy, as sketched in Section 3.2. We begin by specifying the details of the economy's production technology, and then derive our Phillips curve (11).

Technology. A unit continuum of unions, indexed by $k \in [0, 1]$, differentiate labor into distinct tasks. Union k aggregates efficiency units into the union-specific task $\ell_{kt} = \int e_{it}\ell_{ikt}di$, where ℓ_{ikt} are the hours worked supplied by household i to union k. A competitive final goods producer then packages these tasks using the technology

$$y_t = \left(\int_k \ell_{kt}^{\frac{\eta_t - 1}{\eta_t}} dk \right)^{\frac{\eta_t}{\eta_t - 1}}.$$

The price index of a unit of the overall labor aggregate is

$$w_t = \left(\int w_{kt}^{1-\eta_t} dk \right)^{1/(1-\eta_t)},$$

where w_{kt} is the price of the task supplied by union k. Marginal cost pricing by final goods producers requires $p_t = w_t$. The resulting demand for labor from union k is

$$\ell_{kt} = \left(\frac{w_{kt}}{w_t} \right)^{-\eta_t} y_t. \tag{A.1}$$

Integrating both sides across k yields the aggregate production (12), where $d_t \equiv \int \left(\frac{w_{kt}}{w_t} \right)^{-\eta_t} dk$, with ℓ_t denoting total effective hours supplied by households and d_t capturing the efficiency losses due to price dispersion.

From union problem to Phillips curve. We assume that union wage payments to households are subsidized at gross rate $\frac{\bar{\eta}\Xi}{(\bar{\eta}-1)(1-\tau_y)}$, where $\bar{\eta}$ is the steady state elasticity of substitution between varieties of labor and the term Ξ accounts for the fact that the social planner may weight households differently from the labor union. We derive the precise value of Ξ in Appendix D.3; for the purposes of our analysis here, it suffices to note that the labor subsidy takes this general form and that it is financed with a lump-sum tax on unions. The union's problem is to choose the reset wage w^* and ℓ_{kt} to maximize

$$\sum_{s\geq 0}\beta^s\theta^s\left[u_c(c_{t+s})(1-\tau_y)\frac{\bar{\eta}\Xi}{(\bar{\eta}-1)(1-\tau_y)}\frac{w^*}{p_{t+s}}\ell_{kt}-\nu_\ell\left(\ell_{t+s}\right)\ell_{kt}\right]$$

subject to (A.1) and taking c_{t+s} and ℓ_{t+s} as given (since the individual labor union is atomistic). The first-order condition is

$$\sum_{s\geq 0}\beta^s\theta^s\nu_\ell\left(\ell_{t+s}\right)y_{t+s}\eta_{t+s}\left(\frac{p_{t+s}}{p_t}\right)^{\eta_{t+s}}=\frac{\bar{\eta}}{(\bar{\eta}-1)}\sum_{s\geq 0}\beta^s\theta^s\Xi u_c(c_{t+s})(\eta_{t+s}-1)\frac{w_t^*}{p_t}\left(\frac{p_{t+s}}{p_t}\right)^{\eta_{t+s}-1}y_{t+s}, \tag{A.2}$$

where w_t^* is the optimal reset wage chosen at date t. Log-linearizing the first-order condition around a zero-inflation steady state:

$$\sum_{s\geq 0}\beta^s\theta^s\left(\phi\widehat{y}_{t+s}+\widehat{\eta}_{t+s}+\bar{\eta}\left(\widehat{p}_{t+s}-\widehat{p}_t\right)+\widehat{y}_{t+s}-\widehat{y}_{t+s}-\frac{\bar{\eta}}{\bar{\eta}-1}\widehat{\eta}_{t+s}-\widehat{w}_t^*+\bar{\eta}\widehat{p}_t-(\bar{\eta}-1)\widehat{p}_{t+s}+\gamma\widehat{y}_{t+s}\right)=0,$$

where $\phi\equiv\frac{\nu_{\ell\ell}(\bar{\ell})\bar{\ell}}{\nu_\ell(\bar{\ell})}$ and we have used the fact $\widehat{\ell}_t=\widehat{y}_t$ in a first-order approximation of the dynamics. Rearranging

$$\widehat{w}_t^*-\widehat{p}_t=(1-\beta\theta)\sum_{s\geq 0}\beta^s\theta^s\left(\left(\phi+\gamma\right)\widehat{y}_{t+s}-\frac{1}{\bar{\eta}-1}\widehat{\eta}_{t+s}+\widehat{p}_{t+s}-\widehat{p}_t\right)$$

Next, we from the definition of the price index have

$$1+\pi_t\equiv\frac{p_t}{p_{t-1}}=\left(\theta^{-1}-\frac{1-\theta}{\theta}\left(\frac{w_t^*}{p_t}\right)^{1-\eta_t}\right)^{\frac{1}{\eta_t-1}}. \tag{A.3}$$

Log-linearizing around a zero inflation steady state this gives

$$\widehat{\pi}_t=\widehat{p}_t-\widehat{p}_{t-1}=\frac{1-\theta}{\theta}\left(\widehat{w}_t^*-\widehat{p}_t\right).$$

Eliminating $\widehat{w}_t^* - \widehat{p}_t$ and simplifying, we get

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

where $\kappa = \frac{(1-\theta)(1-\beta\theta)(\phi+\gamma)}{\theta}$ and $\psi = -\frac{\kappa}{(\bar{\eta}-1)(\phi+\gamma)}$.

A.2 Werning (2015) special case

This section elaborates on our discussion in Section 5.2 of optimal policy in a special case of our model where monetary policy is distributionally neutral, following Werning (2015). We proceed in three steps. First, we present the assumptions required to arrive at this special case. Second, we derive the distributional irrelevance result. And third, we formally state implications for optimal policy design.

Model assumptions. We consider a special case of our environment, adapted to be to be consistent with the assumptions in Werning (2015).

First, we further restrict the household consumption-savings problem.

Assumption A.1. *Household utility is logarithmic ($\gamma = 1$). The distribution of household productivity e_{it} is acylical (i.e., the mapping Φ from idiosyncratic events to productivity is independent of y_t) and households can self-insure only through saving, not borrowing ($\underline{a} = 0$).*

Second, we assume that the government-supplied asset has particular properties.

Assumption A.2. *The government supplies a perpetuity ($\delta = 0$) whose returns are perfectly indexed to inflation ($\chi_\pi = 1$) and output ($\chi_y = 1$). The supply of the asset is fixed over time ($b_t = \bar{b}$ for all t).*

The third and final assumption restricts the exogenous component of transfers, τ_x.

Assumption A.3. *In steady-state, the exogenous and endogenous components of government transfers to households are zero, i.e. $\bar{\tau}_x = 0$ and $\bar{\tau}_e = 0$.*

Under similar assumptions, Werning (2015) proves that the demand side of the economy responds to monetary policy exactly like the conventional Euler equation (19). Household heterogeneity affects the *split* of the consumption response into indirect income and direct interest rate effects, but leaves the overall *sum* unchanged. We will see that the same logic allows a sharp characterization of optimal policy rules for the Ramsey planner.

The distributional effects of monetary policy. Following steps similar to those in Werning (2015), we can prove the following useful building block result.

Proposition A.1. *Under Assumptions A.1 to A.3, we have that*

$$\Theta_{\omega(\zeta),i} = \mathbf{0} \qquad \forall \zeta. \tag{A.4}$$

In words, changes in nominal interest rates have no effect on the consumption *distribution*, at any horizon.

Optimal policy characterization. Combining Proposition A.1 and our general characterizations of optimal forecast targeting policy rules in Section 2.1, it follows immediately that the presence of the inequality term in the policymaker loss function does not at all affect the target criterion for optimal monetary policy. Corollary A.1 summarizes this conclusion.

Corollary A.1. *Under Assumptions A.1 to A.3 and the conditions of Proposition 1, the optimal monetary policy rule for a Ramsey policymaker with loss function* (29) *can be written as the forecast target criterion*

$$\widehat{\pi}_t + \frac{1}{\bar{\eta}}\left(\widehat{y}_t - \widehat{y}_{t-1}\right) = 0, \quad \forall t = 0, 1, \ldots \tag{A.5}$$

Corollary A.1 formalizes our intuitive discussion in Section 5.2. Finally we also note the following implication of these results: since monetary policy stabilizes output and inflation as well as possible, and since fiscal policy has no additional scope to help with this stabilization (recall Section 4), it follows that fiscal transfer policy is exclusively concerned with inequality stabilization, with the optimal transfer target criterion given as

$$\int \lambda_{\omega(\zeta)} \Theta'_{\omega(\zeta),\tau_x} W \widehat{\boldsymbol{\omega}}(\zeta) d\Gamma(\zeta) = \mathbf{0}. \tag{A.6}$$

Corollary A.2 summarizes this conclusion.

Corollary A.2. *Under the optimal joint fiscal-monetary policy, transfers are set following the target criterion* (A.6)*, minimizing the inequality term in the loss function* (29)*. Monetary policy is set residually to enforce the target criterion* (A.5)*, thus attaining the same paths of aggregate inflation and output as under optimal monetary policy alone.*

We thus achieve a strict separation of monetary and fiscal instruments. Corollary A.2 is related to the results in Bilbiie et al. (2021): while those authors find that transfers can be set to perfectly stabilize inequality between two groups of households, transfers in our HANK model are set to stabilize the general inequality term *as well as possible*. In both cases, given fiscal stabilization of inequality, conventional monetary policy then implements the aggregate allocations familiar from standard representative-agent analysis—in their case because inequality is already perfectly stabilized; in our case because conventional monetary policy cannot help any further with the inequality-related loss.

B Supplementary material for Sections 5.3 and 5.4

This appendix presents supplementary material for our quantitative optimal Ramsey policy analysis. Appendix B.1 begins with further details on model calibration, while Appendix B.2 provides various additional results for our particular business-cycle shock experiments.

B.1 Calibration

Income process. Our process for normal household income is a simplified version of the one that appears in Guvenen et al. (2022), and we use a similar estimation procedure as them in fitting the parameters of the process. We briefly summarize the moments that are targeted here and describe how we deviate from Guvenen et al. while referring the reader to that paper for most of the details.

We estimate the following income process:

$$\begin{aligned}
\log \zeta_{it} &= \rho_\zeta \log \zeta_{it-1} + \xi_{it} \\
\log e_{it} &= \mu_i + \log\left(\zeta_{it}\right)\left(1 + m_t + \alpha \log y_t\right) + \log \bar{e}_t + \xi_{it}^T \\
\xi_{it} &\sim N(0, \sigma_\xi^2) \\
\mu_i &\sim N(0, \sigma_\alpha^2).
\end{aligned}$$

As described by Guvenen et al., over the lifecycle, the cross-sectional variance of earnings grows almost linearly in the age of a cohort. In the absence of dispersion in the deterministic component of individual income profiles, this implies near-random walk behavior in individual earnings. To have a stationary distribution of income, we fix ρ_ζ at a value less than one, setting $\rho_\zeta = 0.9^{1/4}$.[27] ξ_{it}^T is a transitory income shock that follows a two-state Markov chain. We fix the parameters of this Markov chain to mimic the "non-employment" shock in Guvenen et al..[28] Finally μ_i is an individual fixed effect.

When solving the model we discretize the process for ζ_{it}. We furthermore allow for two high-income states, helping the model generate a concentrated distribution of wealth. Workers enter either one of the high-income states with a constant probability, and then

[27] Allowing for heterogeneous income profiles, one finds a lower value of ρ_ζ.

[28] We include this shock in our income process estimation so that the persistent shocks ξ_{it} are not forced to account for all features of the income data. For the sake of parsimony, however, we do not include this shock in quantitative model analysis. We have experimented with including this shock and found the results to be similar; the main difference is that the transitory shock leads to a lower average MPC in steady state.

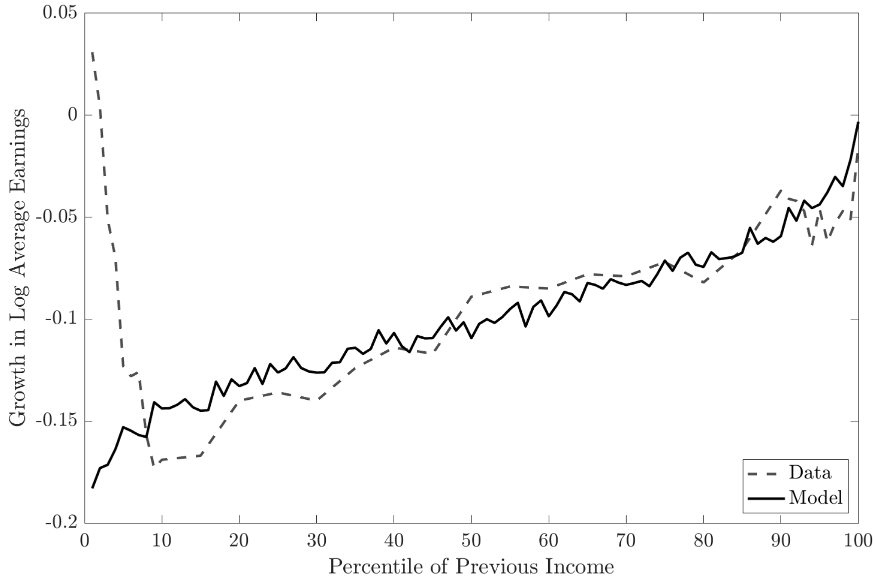

Figure B.1: Incidence of earnings losses during the 1979-1983 recession. The horizontal axis ranks individuals according to their earnings during 1974-1978. For each percentile of this distribution, the vertical axis shows the change from 1979 to 1983 in the log of average earnings for that group. The upward slope implies that lower-income individuals suffered larger earnings losses on average in the recession. Data are taken from Guvenen et al. (2014). The model results are from a simulation of the fitted income process, feeding in the time series of average (aggregate) earnings growth.

later exit with a similarly constant probability. The income levels in these states and the entry and exit probabilities are calibrated to minimize the sum of squared deviations from the moments reported in Table 1. The calibrated process has one state that corresponds to the top 1% of the earnings distribution and another that corresponds to the rest of the top 10%. Households in the top 1% remain there for 26 years in expectation, while households in the rest of the top 10% remain there for 4.4 years on average.[29] The fitted levels of earnings for the high earners are 20 and 3.4 times average earnings, respectively.

There are three parameters that we need to estimate via simulated method of moments estimation: the sensitivity of income dispersion to the business cycle, α; the variance of the persistent shock, σ_ξ^2; and the variance of the fixed effect, σ_μ^2. We simulate the process at a quarterly frequency and time aggregate to annual observations. Many of the moments we target reflect the shape of the earnings growth distribution for earnings growth at 1-year, 3-year, and 5-year horizons. For the purposes of this paper, the main question is how we

[29]The quarterly exit probabilities are 1/104.2 and 1/17.7, respectively.

identify α. Here we proceed as follows. For each business cycle episode between 1979 and 2010, Guvenen et al. (2014) construct a figure analogous to Figure B.1, using average income over the five years prior to the business cycle episode to rank individuals in the distribution. For each recession and expansion, we fit a trend line between the 11th and 80th percentiles of Figure B.1 to obtain a target "slope." We then seek to replicate these slopes in our simulated process. To get the model-implied analogue, we feed in average earnings as our measure of y_t, from there recover household earnings, and then finally construct the model-implied slope. One component of our objective function is then simply the percentage deviation between the model and empirical slopes of the incidence of the business cycle. Overall, the resulting parameters are $\alpha = -1.61$, $\sigma_\xi = 0.077$ and $\sigma_\mu = 0.63$.

Monetary policy shocks for internal calibration. Three parameters of our model are set by internal calibration, to target estimated impulse response functions to identified monetary policy shocks: the relative risk aversion γ; the Phillips curve slope κ; and the output sensitivity of bond returns, χ_y. We use the high-frequency monetary shocks identified by Gertler & Karadi (2015) to estimate the response of key macro aggregates to identified monetary shocks. We first describe the results and internal calibration procedure and then give the details of the empirical implementation.

Figure B.2 shows point estimates for four estimated impulse response functions. In the top-left panel we show the time path of real interest rates, which we construct as $i_t - \pi_{t+1}$. As expected, the expansionary monetary shock leads to a persistent decline in real interest rates. The top-right panel shows that output increases, eventually reaching a peak of 0.8% above steady state. The bottom-left panel reveals that inflation rises quite persistently, with a peak increase of 0.3%. Lastly, the bottom-right panel shows household net worth initially increases by 4%, before then rising a bit further.

Our internal calibration procedure assumes that monetary policy is set according to a standard Taylor rule subject to shocks. We find a path of current and anticipated monetary policy shocks that, when announced at date 0, leads to a change in real rate expectations that perfectly replicates the empirical estimate in Figure B.2.[30] We then choose $\{\gamma, \kappa, \chi_y\}$ to match the peak responses of output and inflation as well as the initial response of household net worth. Figure B.2 shows the model-implied impulse response functions. For output

[30]By the results in McKay & Wolf (2021), the output of this procedure is independent of the baseline policy rule, as long as it induces a determinate, sunspot-free equilibrium. Our choice of baseline rule is thus a purely computational device.

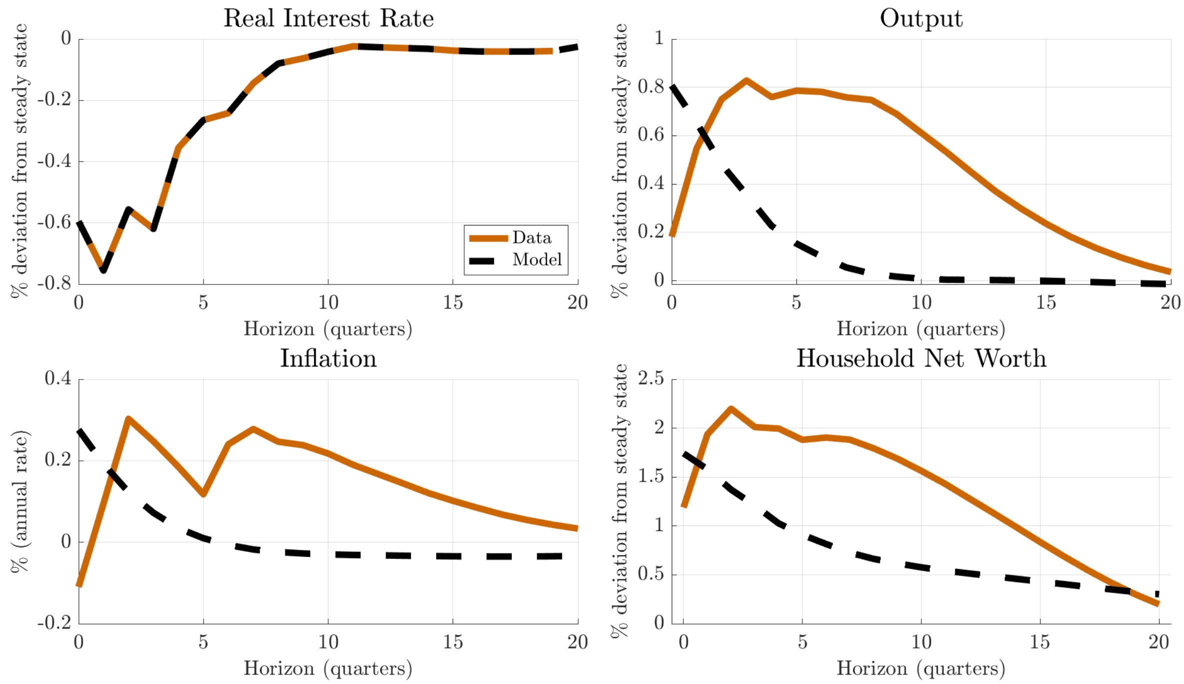

Figure B.2: Empirical (orange) and model-implied (blue-dashed) impulse response following an expansionary monetary policy shock.

and inflation, the model generates the correct peak response, but unsurprisingly does not generate the hump-shaped pattern of the empirical estimates.

Our empirical analysis uses the following series: the 3-month Treasury rate, the inflation rate constructed from the log difference in the GDP deflator, log real GDP per capita, and the log ratio of household net worth from the Financial Accounts of the United States (formerly Flow of Funds) to nominal GDP.[31] Our monetary shock series is equal to the OLS point estimates of monetary shocks as implied by the Gertler & Karadi SVAR-IV, estimated only for scheduled FOMC meetings. We use data from 1980Q3 to 2015Q3, and estimate the impulse response functions using a recursive VAR with the identified shock ordered first (Plagborg-Møller & Wolf, 2021).

[31] In Figure B.2 we plot $W \equiv \log(PW/PY) + \log(Y)$, where PW is nominal net worth, PY is nominal GDP and Y is real GDP.

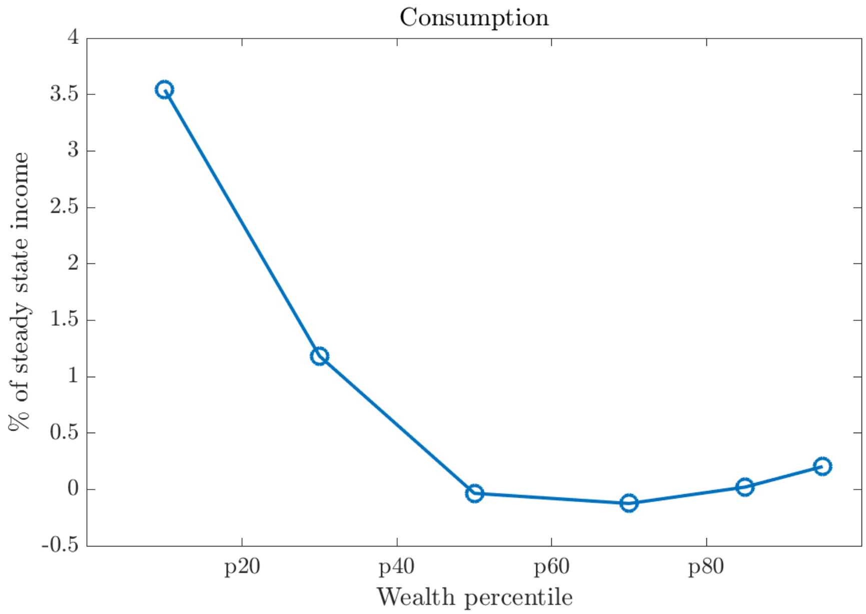

Figure B.3: Initial response of consumption to a \$500 fiscal stimulus payment across the income distribution. The figure includes general equilibrium effects where monetary policy is assumed to follow a simple interest rate rule $i_t = \frac{3}{2}\pi_t + \frac{1}{8}y_t$.

B.2 Additional optimal policy results

We here collect some figures that supplement the quantitative optimal policy analysis presented in Section 5.4.

Distributional effects of stimulus checks. Stimulus checks and interest rate cuts have the same effects on aggregate output and inflation (Wolf, 2021), but they differ substantially in their cross-sectional incidence. Figure 2 in the main text shows the model-implied incidence profile of interest rate cuts; Figure B.3 here does the same for stimulus checks. As expected, the incidence profile here is strongly downward-sloping: the consumption of poor households responds the most, reflecting their high MPCs and low overall income.

Stabilizing consumption at the bottom. What happens if monetary policy is used to stabilize consumption of poor households in the face of our distributional shock m_t? In Figure B.4 we plot the implications of following a monetary policy targeting rule that adds a distributional consideration to the dual mandate targeting rule that seeks to stabilize

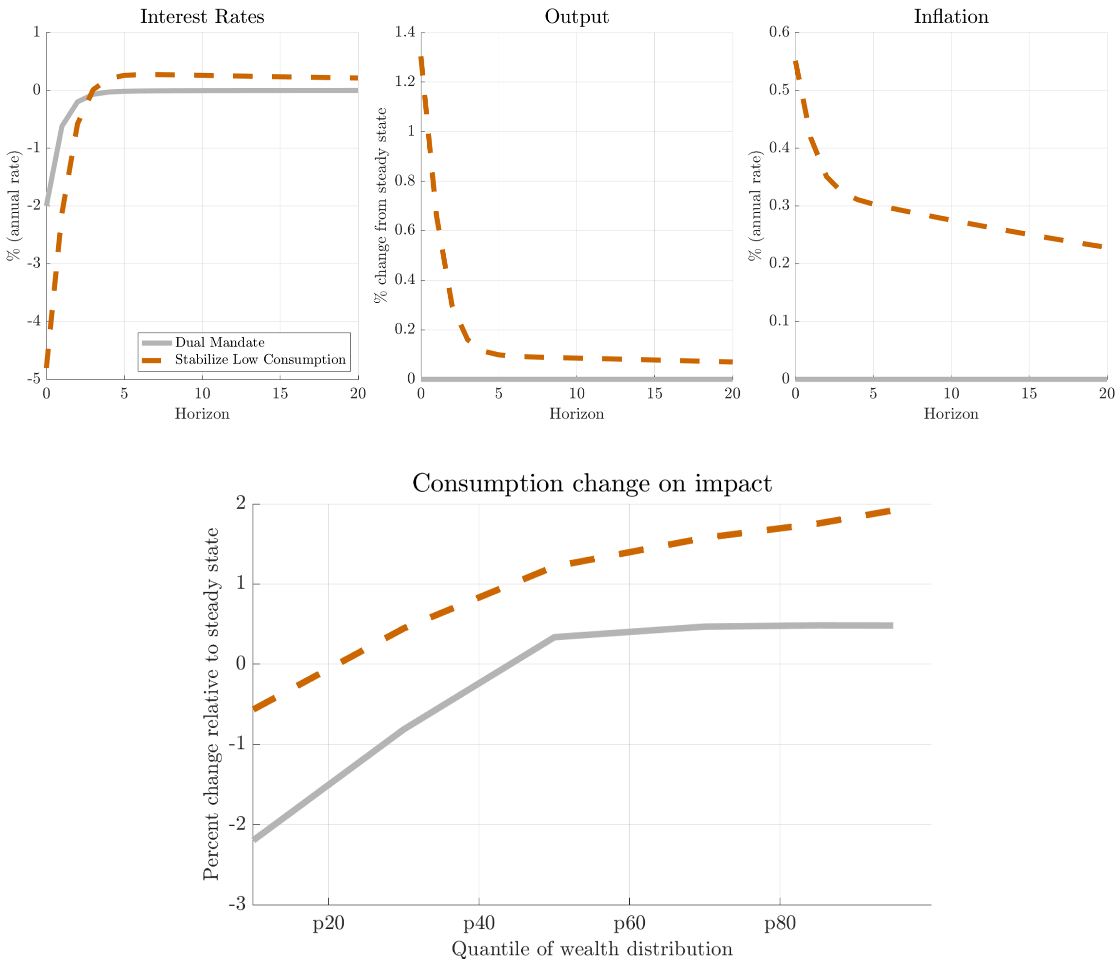

Figure B.4: Simulation of income distribution shock using ad hoc policy rule that stabilizes the consumption of households with low wealth and low income. The bottom panel shows the consumption of wealthy households increases substantially under this rule leading output and inflation to exceed their targets.

the consumption of households with low wealth and low income.[32] As the figure shows, output and inflation exceed the dual-mandate targets and the consumption of high-wealth households is substantially above the steady state level.

[32]The targeting rule is $\hat{\pi}_t + \frac{1}{\bar{\eta}}\left(\hat{y}_t - \hat{y}_{t-1}\right) + \frac{1}{2}\hat{c}_t^{\text{Low}} = 0$, where c_t^{Low} is the aggregated consumption of households with income in the bottom third of the income distribution and wealth below the median.

C Computational appendix

This appendix provides supplementary information on our computational approach. We compute sequence-space transition paths using the methodology developed in Auclert et al. (2021), as discussed further in Appendix C.1. Our computation of the inequality term in the full Ramsey loss function is described in Appendix C.2.

C.1 General equilibrium transition paths

The baseline constraints (2) in our linear-quadratic policy problem in Section 2 are expressed in sequence space. To compute the corresponding constraints (17) - (18) for our "HANK" model, we thus follow the computational techniques of Auclert et al. (2021) to compute the required sequence-space Jacobian matrices. In particular, our computation of the $\mathcal{C}_\bullet$ maps in the augmented HANK "IS curve" (18) leverages the so-called "fake news" algorithm.

For computation of the alternative (but equivalent) constraint formulations (21) - (32) in impulse response space, we follow McKay & Wolf (2021) and proceed as follows. First, we close the model with arbitrary policy rules for the two instruments $\boldsymbol{i}$ and $\boldsymbol{\tau}_x$, subject only to the requirement that the two rules induce a unique equilibrium. We then compute impulse responses of all policy targets to the full menu of contemporaneous and news shocks to those two policy rules. Now denote the impulse response matrix of some variable x_i to shocks to the rule for instrument z_j by $\tilde{\Theta}_{x_i,z_j}$, and similarly write $\tilde{\Theta}_{x_i,\varepsilon_q}$ for responses to non-policy shocks ε_q under the (arbitrary) baseline rule. Finally write $\tilde{\Theta}_{z_j,z_j}$ for the impulse response matrix of the instrument itself. We then define $\Theta_{x_i,z_j} \equiv \tilde{\Theta}_{x_i,z_j}\tilde{\Theta}_{z_j,z_j}^{-1}$ and similarly $\Theta_{x_i,\varepsilon_q} \equiv \tilde{\Theta}_{x_i,\varepsilon_q}\tilde{\Theta}_{z_j,z_j}^{-1}$. The results in McKay & Wolf (2021) imply that the resulting impulse responses are independent of the chosen baseline policy rule.

C.2 Inequality term

The inequality term says that the planner would like to stabilize a very large number of targets—one for each history ζ. Both for intuition and for computation, it is useful to observe that these consumption shares will only fluctuate if the inputs to the household's decision problem fluctuate. By our discussion of the consumption-savings problem in Section 3.1, those inputs include total labor income (equal to $\boldsymbol{y}$), the return on bonds ($\boldsymbol{r}$), transfers ($\boldsymbol{\tau}_x$ and $\boldsymbol{\tau}_e$), and the inequality shock ($\boldsymbol{m}$). In this section we leverage this insight to recast the problem of stabilizing consumption shares as one of stabilizing the inputs to the household consumption-savings decision, thus giving the representation in (33), couched in terms of a small number of aggregates rather than a distribution of histories.

Reformulating the inequality term. Let $\boldsymbol{x} \equiv (\boldsymbol{r}', \boldsymbol{y}', \boldsymbol{\tau}_x', \boldsymbol{\tau}_e', \boldsymbol{m}')'$ be the stacked sequences of inputs to the household problem. Our goal is to show that there is symmetric matrix Q such that

$$\sum_{t=0}^{\infty} \beta^t \int \frac{\widehat{\omega}_t(\zeta, \boldsymbol{x})^2}{\bar{\omega}(\zeta)} d\Gamma(\zeta) = \widehat{\boldsymbol{x}}' Q \widehat{\boldsymbol{x}} + \mathcal{O}(||\widehat{\boldsymbol{x}}||^3),$$

where here we have been explicit that the consumption shares at date t depend on the full sequences of inputs $\boldsymbol{x}$. To arrive at this representation, consider a first-order approximation to the time-t consumption share of individuals with history ζ:

$$\widehat{\omega}_t(\zeta, \boldsymbol{x}) \approx \Omega_t(\zeta) \widehat{\boldsymbol{x}}$$

where the derivative $\Omega_t(\zeta)$ will be defined formally below. This yields

$$\frac{\widehat{\omega}_t(\zeta^t, \boldsymbol{x})^2}{\bar{\omega}(\zeta^t)} = \widehat{\boldsymbol{x}}' \underbrace{\frac{\Omega_t(\zeta^t)' \Omega_t(\zeta^t)}{\bar{\omega}(\zeta^t)}}_{\equiv Q_t(\zeta^t)} \widehat{\boldsymbol{x}} + \mathcal{O}(||\widehat{\boldsymbol{x}}||^3).$$

We then integrate across histories and take the discounted sum across time to arrive at

$$\sum_{t=0}^{\infty} \beta^t \int \frac{\widehat{\omega}_t(\zeta^t, \boldsymbol{x})^2}{\bar{\omega}(\zeta^t)} d\Gamma(\zeta^t) = \widehat{\boldsymbol{x}}' \underbrace{\left(\sum_{t=0}^{\infty} \beta^t \int Q_t(\zeta^t) d\Gamma(\zeta^t) \right)}_{\equiv Q} \widehat{\boldsymbol{x}} + \mathcal{O}(||\widehat{\boldsymbol{x}}||^3),$$

giving the desired representation. We have thus arrived at a representation that almost fits into our general linear-quadratic set-up of Section 2.1, the sole difference being that the objective function features a non-diagonal quadratic form. To extend our optimal policy results to this more general case, we consider the same problem as in Section 2.1, but replacing the diagonal loss function (1) by the more general (non-diagonal) expression

$$\mathcal{L} \equiv \frac{1}{2} \boldsymbol{x}' P \boldsymbol{x}, \tag{C.1}$$

The corresponding necessary and sufficient first-order conditions yield the more general optimal policy targeting rule

$$\Theta_{xz}' P \boldsymbol{x} = 0$$

and the optimal instrument path

$$\boldsymbol{z}^* \;\equiv\; -\left(\Theta'_{x,z} P \Theta_{x,z}\right)^{-1} \times \left(\Theta'_{x,z} P \Theta_{x,\varepsilon} \cdot \boldsymbol{\varepsilon}\right)$$

Evolution of consumption shares. We now explain the first-order approximation of consumption shares

$$\widehat{\omega}_t(\zeta^t) \approx \Omega_t(\zeta^t)\boldsymbol{x}$$

that we used above. Let $c_t(\zeta^t, \boldsymbol{x})$ be the consumption in date t after history ζ^t with the input sequences given by $\boldsymbol{x}$, and similarly let $a_t(\zeta^t, \boldsymbol{x})$ be the savings chosen in date t. Also let ζ_t be the date-t value of the idiosyncratic state. Using the standard recursive representation of the household's problem, we can write these choices in terms policy functions, f and g, that take as their arguments assets $a_{t-1}(\zeta^{t-1})$ and the current shock ζ_t, so we have

$$c_t(\zeta^t, \boldsymbol{x}) \;=\; f_t\left(a_{t-1}(\zeta^{t-1}, \boldsymbol{x}), \zeta_t, \boldsymbol{x}\right) \tag{C.2}$$

$$a_t(\zeta^t, \boldsymbol{x}) \;=\; g_t\left(a_{t-1}(\zeta^{t-1}, \boldsymbol{x}), \zeta_t, \boldsymbol{x}\right). \tag{C.3}$$

We now consider a first-order approximation to f and g around $\boldsymbol{x} = \bar{\boldsymbol{x}}$:

$$\begin{aligned} c_t(\zeta^t, \boldsymbol{x}) &\approx \bar{c}(\zeta^t) + \frac{dc_t(\zeta^t, \boldsymbol{x})}{d\boldsymbol{x}}\widehat{\boldsymbol{x}} \\ a_t(\zeta^t, \boldsymbol{x}) &\approx \bar{a}(\zeta^t) + \frac{da_t(\zeta^t, \bar{\boldsymbol{x}})}{d\boldsymbol{x}}\widehat{\boldsymbol{x}}. \end{aligned}$$

The derivatives that appear here are total derivatives with respect to $\boldsymbol{x}$, including both the effect on the policy rule at date t and the effect on assets $a_{t-1}(\zeta^{t-1}, \boldsymbol{x})$. The derivatives are evaluated at the steady-state inputs $\bar{\boldsymbol{x}}$ and the level of assets that an individual with history ζ^t would have if the inputs x remained at steady state forever, which we denote by $\bar{a}(\zeta^{t-1})$. To calculate these derivatives, we differentiate (C.2) and (C.3):

$$\frac{dc_t(\zeta^t, \bar{\boldsymbol{x}})}{d\boldsymbol{x}} = \frac{\partial f_t\left(\bar{a}_{t-1}(\zeta^{t-1}), \zeta_t, \bar{\boldsymbol{x}}\right)}{\partial a}\frac{da_{t-1}(\zeta^{t-1}, \bar{\boldsymbol{x}})}{d\boldsymbol{x}} + \frac{\partial f_t\left(\bar{a}_{t-1}(\zeta^{t-1}), \zeta_t, \bar{\boldsymbol{x}}\right)}{\partial \boldsymbol{x}} \tag{C.4}$$

$$\frac{da_t(\zeta^t, \bar{\boldsymbol{x}})}{d\boldsymbol{x}} = \frac{\partial g_t\left(\bar{a}_{t-1}(\zeta^{t-1}), \zeta_t, \bar{\boldsymbol{x}}\right)}{\partial a}\frac{da_{t-1}(\zeta^{t-1}, \bar{\boldsymbol{x}})}{d\boldsymbol{x}} + \frac{\partial g_t\left(\bar{a}_{t-1}(\zeta^{t-1}), \zeta_t, \bar{\boldsymbol{x}}\right)}{\partial \boldsymbol{x}}. \tag{C.5}$$

The partial derivative $\frac{\partial f_t\left(\bar{a}_{t-1}(\zeta^{t-1}), \zeta_t, \bar{\boldsymbol{x}}\right)}{\partial a}$ is the marginal propensity to consume for an individual with states $(\bar{a}_{t-1}(\zeta^{t-1}), \zeta_t)$ in the stationary equilibrium, and $\frac{\partial g_t\left(\bar{a}_{t-1}(\zeta^{t-1}), \zeta_t, \bar{\boldsymbol{x}}\right)}{\partial a}$ is the marginal propensity to save. Similarly, the partial derivative $\frac{\partial f_t\left(\bar{a}_t(\zeta^{t-1}), \zeta_t, \bar{\boldsymbol{x}}\right)}{\partial \boldsymbol{x}}$ is the derivative of the consumption policy rule with respect to the input sequences for an individual with the same states, and $\frac{\partial g_t\left(\bar{a}_t(\zeta^{t-1}),\zeta_t,\bar{\boldsymbol{x}}\right)}{\partial \boldsymbol{x}}$ is the analogous derivative of the savings policy rule.

We discuss below how to compute these derivatives. Given that they have been recovered, it remains to move from consumption *levels* to consumption *shares*:

$$\omega_t(\zeta^t, \boldsymbol{x}) \equiv \frac{c_t(\zeta^t, \boldsymbol{x})}{\int c_t(\zeta^t, \boldsymbol{x}) d\Gamma_t(\zeta)} \approx \bar{\omega}_t(\zeta^t) + \frac{1}{\bar{c}} \frac{dc_t(\zeta^t, \boldsymbol{x})}{d\boldsymbol{x}} \widehat{\boldsymbol{x}} - \int \frac{\bar{c}(\zeta^t)}{\bar{c}^2} \frac{dc_t(\zeta^t, \boldsymbol{x})}{d\boldsymbol{x}} d\Gamma(\zeta^t) \widehat{\boldsymbol{x}}.$$

$\Omega_t(\zeta^t)$ is then given by

$$\Omega_t(\zeta^t) \equiv \frac{1}{\bar{c}} \frac{dc_t(\zeta^t, \boldsymbol{x})}{d\boldsymbol{x}} - \int \frac{\bar{c}(\zeta^t)}{\bar{c}^2} \frac{dc_t(\zeta^t, \boldsymbol{x})}{d\boldsymbol{x}} d\Gamma(\zeta^t). \tag{C.6}$$

Computing Q. To compute $\Omega_t(\zeta^t)$ and so Q, the key challenge is to arrive at the derivatives in (C.2) - (C.3). To do so we begin by simulating a history ζ^t for $t = 0, 1, ..., T$ in a stationary equilibrium (i.e. with $\boldsymbol{x} = \bar{\boldsymbol{x}}$). At each date along this simulation, we recover the required partial derivatives as follows. The marginal propensities to consume and save can be computed by standard methods. For the derivatives of the policy rules, we use the fact that the derivatives with respect to past prices are zero and the derivatives with respect to current and future prices only depends on the number of periods until the price change occurs. This allows us to compute all the derivatives by perturbing prices at a single date and iterating backwards in time using a single loop from T to 0 (see Auclert et al., 2021). With the partial derivatives in hand, we then construct $\frac{dc_t(\zeta^t,\bar{\boldsymbol{x}})}{d\boldsymbol{x}}$ and $\frac{da_t(\zeta^t,\bar{\boldsymbol{x}})}{d\boldsymbol{x}}$ by iterating (C.4)-(C.5) forward starting with $\frac{da_{-1}(\zeta^{-1},\bar{\boldsymbol{x}})}{d\boldsymbol{x}} = 0$. This initial condition reflects the fact that assets (before interest) entering date 0 are pre-determined with respect to the prices in $\boldsymbol{x}$ that apply from date 0 onwards.

Given those derivatives, we can recover $\Omega_t(\zeta)$ and thus get $Q_t(\zeta^t)$ as well as Q itself.

Impulse response representation. Using (C.4), (C.5), and (C.6), we can write $\omega_t(\zeta)$ as a linear function of the aggregate variables contained in $\boldsymbol{x}$. Since each $\omega_t(\zeta)$ is linearly related to $\boldsymbol{x}$, and since we can recover the entries of $\boldsymbol{x}$ as a function of $(\boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{i}, \boldsymbol{\tau}_x, \boldsymbol{m})$ (by Lemma 1), we recover the impulse response representation in (31).

D Proofs and auxiliary lemmas

D.1 Proof of Lemma 1

We begin the proof by re-stating and slightly simplifying the definition of an equilibrium in Definition 1. We first repeat the Phillips curve in stacked form:

$$\Pi_\pi \widehat{\boldsymbol{\pi}} = \Pi_y \widehat{\boldsymbol{y}} + \psi \boldsymbol{\varepsilon}, \tag{D.1}$$

where

$$\Pi_\pi = \begin{pmatrix} 1 & -\beta & 0 & \cdots \\ 0 & 1 & -\beta & \cdots \\ 0 & 0 & 1 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}, \quad \Pi_y = \kappa I.$$

Turning to the demand side, we re-write (10) as

$$\widehat{\boldsymbol{y}} = \mathcal{C}_y \widehat{\boldsymbol{y}} + \mathcal{C}_r \widehat{\boldsymbol{r}} + \mathcal{C}_x \widehat{\boldsymbol{\tau}}_x + \mathcal{C}_e \widehat{\boldsymbol{\tau}}_e + \mathcal{C}_m \widehat{\boldsymbol{m}}, \tag{D.2}$$

where we have used the equilibrium relationships $\widehat{c}_t = \widehat{y}_t$, $w_t/p_t = 1$ and $\widehat{y}_t = \widehat{\ell}_t$, and write $\mathcal{C}_y = \mathcal{C}_\ell = \mathcal{C}_{w/p}$. Next, using (13) and (14), we write the relationships between asset prices and rates of return as

$$\widehat{r}_0 = r_0(\widehat{\pi}_0, \widehat{y}_0, \widehat{q}_0) \tag{D.3}$$

$$\widehat{\boldsymbol{r}}_{+1} = r_{+1}(\widehat{\boldsymbol{i}}, \widehat{\boldsymbol{\pi}}) \tag{D.4}$$

and

$$\widehat{\boldsymbol{q}} = q(\widehat{\boldsymbol{\pi}}_{+1}, \widehat{\boldsymbol{y}}_{+1}, \widehat{\boldsymbol{r}}_{+1}) \tag{D.5}$$

Finally, we combine the government budget constraint (15) and the law of motion for government debt (16) and solve for $\tau_{e,t}$ to obtain

$$\widehat{\boldsymbol{\tau}}_e = \tau_e(\widehat{\boldsymbol{y}}, \widehat{\boldsymbol{\tau}}_x, \widehat{\boldsymbol{\pi}}, \widehat{\boldsymbol{q}}). \tag{D.6}$$

Our first auxiliary result is that, given the shocks $(\boldsymbol{m}, \boldsymbol{\varepsilon})$ and policy choices $(\boldsymbol{i}, \boldsymbol{\tau}_x)$, a list $(\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\tau}_e, \boldsymbol{\pi}, \boldsymbol{q})$ is part of an equilibrium if and only if (D.1) - (D.6) hold. To show this we need to check the conditions of Definition 1: requirements 1 - 3 and goods market clearing hold by the construction of (D.2); requirement 4 is re-stated in (D.1); requirement 5 is imposed by (D.6); requirement 6 is imposed by (D.3)-(D.5); and finally, bond market clearing holds by Walras' Law.

We now simplify this characterization of equilibria further to arrive at Lemma 1. The key step in the argument is to solve out the asset-pricing and government financing rules and plug them into (D.2). First, given $\widehat{\boldsymbol{\pi}}$ and $\widehat{\boldsymbol{i}}$, we can recover $\widehat{\boldsymbol{r}}_{+1}$ from (D.4). We can thus recover $\widehat{\boldsymbol{q}}$ from (D.5), and so finally we get $\widehat{r}_0$ from (D.3). Second, given $\widehat{\boldsymbol{y}}$, $\widehat{\boldsymbol{\tau}}_x$, $\widehat{\boldsymbol{\pi}}$ and $\widehat{\boldsymbol{q}}$, we can recover $\widehat{\boldsymbol{\tau}}_e$ from (D.6). We can thus write

$$\widehat{\boldsymbol{y}} = \mathcal{C}_y\widehat{\boldsymbol{y}} + \mathcal{C}_r\widehat{\boldsymbol{r}}(\widehat{\boldsymbol{y}}, \widehat{\boldsymbol{\pi}}, \widehat{\boldsymbol{i}}) + \mathcal{C}_x\widehat{\boldsymbol{\tau}}_x + \mathcal{C}_e\widehat{\boldsymbol{\tau}}_e(\widehat{\boldsymbol{y}}, \widehat{\boldsymbol{\pi}}, \widehat{\boldsymbol{i}}, \widehat{\boldsymbol{\tau}}_x) + \mathcal{C}_m\boldsymbol{m}$$

and so

$$\widehat{\boldsymbol{y}} = \underbrace{[\mathcal{C}_y + \mathcal{C}_r\mathcal{R}_y + \mathcal{C}_e\mathcal{T}_y]}_{\tilde{\mathcal{C}}_y}\widehat{\boldsymbol{y}} + \underbrace{[\mathcal{C}_r\mathcal{R}_\pi + \mathcal{C}_e\mathcal{T}_\pi]}_{\tilde{\mathcal{C}}_\pi}\widehat{\boldsymbol{\pi}} + \underbrace{[\mathcal{C}_r\mathcal{R}_i + \mathcal{C}_e\mathcal{T}_i]}_{\tilde{\mathcal{C}}_i}\widehat{\boldsymbol{i}} + \underbrace{[\mathcal{C}_x + \mathcal{C}_e\mathcal{T}_x]}_{\tilde{\mathcal{C}}_x}\widehat{\boldsymbol{\tau}}_x + \mathcal{C}_m\boldsymbol{m} \quad \text{(D.7)}$$

where $\mathcal{T}_\bullet$ and $\mathcal{R}_\bullet$ are derivative matrices for the maps $\widehat{\boldsymbol{r}}(\bullet)$ and $\widehat{\boldsymbol{\tau}}_e(\bullet)$. (D.7) embeds (13), (14) and (D.6). We have thus reduced the equilibrium characterization from statements about $(\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\tau}_e, \boldsymbol{\pi}, \boldsymbol{q})$ to statements about $(\boldsymbol{y}, \boldsymbol{\pi})$, establishing the claim. □

D.2 Proof of Proposition 1

In light of Lemma 1, we can re-state the optimal policy problem as minimizing (20) subject to the two constraints (D.1) and (D.7). This problem gives the following necessary and sufficient first-order conditions:

$$\lambda_\pi W\widehat{\boldsymbol{\pi}} + \Pi_\pi' W\boldsymbol{\varphi}_\pi - \tilde{\mathcal{C}}_\pi' W\boldsymbol{\varphi}_y = \boldsymbol{0} \quad \text{(D.8)}$$

$$\lambda_y W\widehat{\boldsymbol{y}} - \Pi_y' W\boldsymbol{\varphi}_\pi + (I - \tilde{\mathcal{C}}_y') W\boldsymbol{\varphi}_y = \boldsymbol{0} \quad \text{(D.9)}$$

$$-\tilde{\mathcal{C}}_i' W\boldsymbol{\varphi}_y = \boldsymbol{0}, \quad \text{(D.10)}$$

where $\boldsymbol{\varphi}_\pi$ and $\boldsymbol{\varphi}_y$ are sequences of Lagrange multipliers on the two constraints. The proof of Proposition 1 proceeds by guessing (and then verifying) that $\boldsymbol{\varphi}_y = \boldsymbol{0}$. Under this assumption, we can combine (D.8) - (D.9) to get

$$\lambda_\pi\widehat{\boldsymbol{\pi}} + \lambda_y W^{-1}\Pi_\pi'(\Pi_y')^{-1}W\widehat{\boldsymbol{y}} = \boldsymbol{0}$$

It is straightforward to verify that this can be re-written as

$$\lambda_\pi \widehat{\boldsymbol{\pi}} + \frac{\lambda_y}{\kappa} \begin{pmatrix} 1 & 0 & 0 & \dots \\ -1 & 1 & 0 & \dots \\ 0 & -1 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \widehat{\boldsymbol{y}} = \mathbf{0} \tag{D.11}$$

But this is just (24), with the conclusion following for $t = 0$ since $\widehat{y}_{-1} = 0$ (as the economy starts from steady state). It now remains to verify the guess that $\boldsymbol{\varphi}_y = \mathbf{0}$. For this, consider some arbitrary $(\boldsymbol{m}, \boldsymbol{\varepsilon})$, and let $(\widehat{\boldsymbol{y}}^*, \widehat{\boldsymbol{\pi}}^*)$ denote the solution of the system (D.1) and (D.11) given $(\boldsymbol{m}, \boldsymbol{\varepsilon})$. Plugging into (D.7) and re-arranging:

$$\underbrace{\widehat{\boldsymbol{y}}^* - \tilde{\mathcal{C}}_y \widehat{\boldsymbol{y}}^* - \tilde{\mathcal{C}}_\pi \widehat{\boldsymbol{\pi}}^* - \mathcal{C}_m \boldsymbol{m}}_{\text{demand target}} = \tilde{\mathcal{C}}_i \widehat{\boldsymbol{i}} \tag{D.12}$$

It thus remains to show that condition (23) is precisely sufficient to ensure that we can always find $\widehat{\boldsymbol{i}}$ such that (D.12) holds. To see this, note that the left-hand side of (D.12) is an excess demand term: supply $\widehat{\boldsymbol{y}}^*$ vs. demand $\tilde{\mathcal{C}}_y \widehat{\boldsymbol{y}}^* + \tilde{\mathcal{C}}_\pi \widehat{\boldsymbol{\pi}}^* + \mathcal{C}_m \boldsymbol{m}$. The fact that the two terms have the same net present value follows from the integrated household and government budget constraint. To see this formally, note that the supply term has net present value

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}}\right)^t \bar{y}\widehat{y}_t \tag{D.13}$$

For the demand term, aggregation of the household budget constraint across all households gives

$$c_t + a_t = (1 + r_t)a_{t-1} + (1 - \tau_y)y_t + \tau_{xt} + \tau_{et}$$

and so

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}}\right)^t \bar{c}\widehat{c}_t = \sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}}\right)^t \{(1+\bar{r})\bar{a}\widehat{r}_t + (1-\tau_y)\bar{y}\widehat{y}_t + \bar{\tau}_x\widehat{\tau}_{xt} + \bar{\tau}_e\widehat{\tau}_{et}\} \tag{D.14}$$

Doing the same for the government budget constraint, we get

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}}\right)^t \{(1+\bar{r})\bar{a}\widehat{r}_t + \bar{\tau}_x\widehat{\tau}_{xt} + \bar{\tau}_e\widehat{\tau}_{et}\} = \sum_{t=0}^{\infty} \tau_y \bar{y}\widehat{y}_t \tag{D.15}$$

Combining (D.15) and (D.14), we get (D.13), as claimed. This completes the argument. □

D.3 Proof of Proposition 2

Let U_t denote the time-t flow utility of the Ramsey planner. To derive the second-order approximation to the social welfare function, it is convenient to begin by writing U_t in terms of log deviations of c_t and ℓ_t from steady state:

$$U_t = \int \varphi(\zeta) \frac{\left(\bar{c}e^{\widehat{c}_t}\omega_t(\zeta)\right)^{1-\gamma} - 1}{1-\gamma} d\Gamma(\zeta) - \nu\left(\bar{\ell}e^{\widehat{\ell}_t}\right). \tag{D.16}$$

Our objective is to construct a second-order approximation of (D.16). Similar to the analysis in Woodford (2003), our strategy is to consider an efficient steady state, allowing evaluation of Equation (D.16) to second order using only a first-order approximation of aggregate equilibrium dynamics.

Optimality of the steady state requires that the weighted marginal utility of consumption is equalized across histories:

$$\varphi(\zeta)\bar{c}^{1-\gamma}\bar{\omega}(\zeta)^{-\gamma} = \bar{u}_c\bar{c} \qquad \forall\zeta$$

for some constant $\bar{u}_c$. Rearranging, we can write this as

$$\varphi(\zeta)^{1/\gamma} = \bar{c}\bar{\omega}(\zeta)\bar{u}_c^{1/\gamma} \qquad \forall\zeta$$

Furthermore imposing that consumption shares integrate to 1 yields

$$\int \varphi(\zeta)^{1/\gamma} d\Gamma(\zeta) = \bar{c}\bar{u}_c^{1/\gamma}. \tag{D.17}$$

Combining the previous two equations, we can recover consumption shares as a function of planner weights:

$$\bar{\omega}(\zeta) = \frac{\varphi(\zeta)^{1/\gamma}}{\int \varphi(\zeta)^{1/\gamma} d\Gamma(\zeta)} \qquad \forall\zeta$$

For future reference it will furthermore be useful to define

$$\Xi \equiv \left(\int \varphi(\zeta)^{1/\gamma} d\Gamma(\zeta)\right)^{\gamma} = \varphi(\zeta)\bar{\omega}(\zeta)^{-\gamma} \qquad \forall\zeta. \tag{D.18}$$

With these preliminary definitions out of the way, we can begin constructing the second-order approximation of (D.16). Differentiating U_t with respect to $\widehat{c}_t$, we find that

$$\begin{aligned}\frac{\partial U}{\partial \widehat{c}_t} &= \int \varphi(\zeta)(\bar{c}\bar{\omega}(\zeta))^{1-\gamma} d\Gamma(\zeta) \\ &= \bar{c}^{1-\gamma}\Xi\end{aligned}$$

where the second line follows from the definition of Ξ and some algebra. Notice that the definition of Ξ and (D.17) together imply that $\Xi = \bar{u}_c/\bar{c}^{-\gamma}$ so Ξ is the ratio of the (common) marginal utility of consumption as evaluated by the planner and the marginal utility of aggregate consumption used by the labor union to value income gains.

Next we have

$$\frac{\partial U}{\partial \widehat{\ell}_t} = -\nu_\ell(\bar{\ell})\bar{\ell}.$$

It follows from the steady state version of equation (A.2) that $\Xi\bar{c}^{-\gamma} = \nu_\ell$. As the steady-state resource constraint is $\bar{c} = \bar{y} = \bar{\ell}$, it follows that $\frac{\partial U}{\partial \widehat{c}_t} + \frac{\partial U}{\partial \widehat{\ell}_t} = 0$, corresponding to efficiency of the total level of aggregate economic activity.

Differentiating U_t with respect to the consumption shares $\omega_t(\zeta)$ we have

$$\begin{aligned}\frac{\partial U}{\partial \omega_t(\zeta)} &= \varphi(\zeta)\bar{c}^{1-\gamma}\bar{\omega}(\zeta)^{-\gamma} d\Gamma(\zeta) \\ &= \bar{c}^{1-\gamma}\Xi d\Gamma(\zeta)\end{aligned}$$

Turning to second order terms, we begin again with the total level and cross-sectional split of consumption. We find

$$\begin{aligned}\frac{\partial^2 U_t}{\partial \widehat{c}_t^2} &= (1-\gamma)\Xi\bar{c}^{1-\gamma} \\ \frac{\partial U_t}{\partial \omega_t(\zeta)^2} &= -\gamma\bar{c}^{1-\gamma}\frac{\Xi}{\bar{\omega}(\zeta)} d\Gamma(\zeta) \\ \frac{\partial^2 U_t}{\partial \widehat{c}_t \partial \omega_t(\zeta)} &= (1-\gamma)\Xi\bar{c}^{1-\gamma} d\Gamma(\zeta)\end{aligned}$$

For hours worked we have

$$\frac{\partial^2 U}{\partial \widehat{\ell}_t^2} = -\nu_{\ell\ell}(\bar{\ell})\bar{\ell}^2 - \nu_\ell(\bar{\ell})\bar{\ell}$$

We can now put everything together, giving the following second-order approximation of time-t planner utility (D.16):

$$\begin{aligned} U_t &\approx \bar{U} + \bar{c}^{1-\gamma}\Xi\widehat{c}_t - \nu_\ell(\bar{\ell})\bar{\ell}\widehat{\ell}_t \\ &+\frac{1}{2}(1-\gamma)\Xi\bar{c}^{1-\gamma}\widehat{c}_t^2 - \frac{1}{2}\left[\nu_{\ell\ell}(\bar{\ell})\bar{\ell}^2 + \nu_\ell(\bar{\ell})\bar{\ell}\right]\widehat{\ell}_t^2 - \frac{1}{2}\gamma\bar{c}^{1-\gamma}\Xi\int\frac{\widehat{\omega}(\zeta)^2}{\bar{\omega}(\zeta)}d\Gamma(\zeta) \\ &+\bar{c}^{1-\gamma}\Xi\int\widehat{\omega}_t(\zeta)d\Gamma(\zeta) + (1-\gamma)\bar{c}^{1-\gamma}\Xi\widehat{c}_t\int\widehat{\omega}_t(\zeta)d\Gamma(\zeta) \end{aligned}$$

Since consumption shares integrate to 1, it follows that $\int\widehat{\omega}_t(\zeta)d\Gamma(\zeta) = 0$, and so all terms in the last row are zero. We now wish to evaluate the remaining terms to second order. To begin, note that the resource constraint and production function give

$$\widehat{c}_t = \widehat{y}_t = \widehat{\ell}_t - \widehat{d}_t$$

where the last term reflects the efficiency loss from wage dispersion. Substituting this in for $\widehat{\ell}_t$ everywhere we have

$$\begin{aligned} U_t &\approx \bar{U} + \bar{c}^{1-\gamma}\Xi\widehat{c}_t - \nu_\ell(\bar{\ell})\bar{\ell}\left(\widehat{c}_t + \widehat{d}_t\right) \\ &+\frac{1}{2}(1-\gamma)\Xi\bar{c}^{1-\gamma}\widehat{c}_t^2 - \frac{1}{2}(\phi+1)\,\nu_\ell(\bar{\ell})\bar{\ell}(\widehat{c}_t + \widehat{d}_t)^2 - \frac{1}{2}\gamma\bar{c}^{1-\gamma}\Xi\int\frac{\widehat{\omega}(\zeta)^2}{\bar{\omega}(\zeta)}d\Gamma(\zeta) \end{aligned}$$

where we have used the definition of ϕ. To simplify this expression further, impose the aggregate resource constraint $\widehat{c}_t = \widehat{y}_t$, use that $\bar{c}^{1-\gamma}\Xi = \nu_\ell(\bar{\ell})\bar{\ell}$, and finally note that all higher-order price dispersion terms can be ignored to second order. We thus get

$$U_t \approx \bar{U} - \nu_\ell(\bar{\ell})\bar{\ell}\widehat{d}_t - \frac{1}{2}\nu_\ell(\bar{\ell})\bar{\ell}(\gamma+\phi)\widehat{y}_t^2 - \frac{1}{2}\gamma\nu_\ell(\bar{\ell})\bar{\ell}\int\frac{\widehat{\omega}(\zeta)^2}{\bar{\omega}(\zeta)}d\Gamma(\zeta)$$

The last step in the derivation is to express $\widehat{d}_t$ in terms of the history of inflation, closely following the arguments in Woodford (2003). Recall that the dispersion term is defined as

$$\begin{aligned} d_t &\equiv \int\left(\frac{w_{kt}}{w_t}\right)^{-\eta_t}dk \\ &= \int\left(\frac{e^{\widehat{w}_{kt}}}{w_t}\right)^{-\eta_t}dk, \end{aligned}$$

where we have defined $\widehat{w}_{kt}$ as the log of w_{kt}. Taking a second-order approximation around

$\widehat{w}_{kt} = \bar{w}_t \equiv \mathbb{E}_k \left[\log w_{kt}\right]$ and $\varepsilon_t = \bar{\eta}$ yields

$$\begin{aligned}\widehat{d}_t &\approx \int -\varepsilon(\widehat{w}_{kt} - \bar{w}_t) + \frac{1}{2}\left[\bar{\eta}^2(\widehat{w}_{kt} - \bar{w}_t)^2 - 2(\widehat{w}_{kt} - \bar{w}_t)\widehat{\varepsilon}_t\right] dk \\ &= \frac{\bar{\eta}^2}{2}\mathrm{Var}_k\left[\widehat{w}_{kt}\right]\end{aligned}$$

where we have simplified using that fact that, at our expansion point, there is no dispersion in w_{kt}, so $e^{\widehat{w}_{kt}} = \bar{w}_t \forall k$. Next we use the Calvo structure to rewrite the definition of d_t as

$$d_t = \theta \int \left(\frac{w_{kt-1}}{w_{t-1}}\right)^{-\eta_t} dk\,(1+\pi_t)^{\eta_t} + (1-\theta)\left(\frac{1}{1-\theta} - \frac{\theta}{1-\theta}(1+\pi_t)^{\varepsilon-1}\right)^{\varepsilon/(\varepsilon-1)}$$

A second-order approximation of this expression (around a zero-inflation steady state) yields

$$\begin{aligned}\widehat{d}_t &\approx \theta\frac{\bar{\varepsilon}^2}{2}\mathrm{Var}_k\left[\widehat{w}_{kt-1}\right] + \frac{\theta\bar{\eta}}{2(1-\theta)}\widehat{\pi}_t^2 \\ &\approx \theta\widehat{d}_{t-1} + \frac{\theta\bar{\eta}}{2(1-\theta)}\widehat{\pi}_t^2.\end{aligned}$$

Solving backwards:

$$\widehat{d}_t \approx \theta^{t+1}\widehat{d}_{-1} + \frac{\theta\bar{\eta}}{2(1-\theta)}\sum_{s=0}^{t}\theta^{t-s}\widehat{\pi}_s^2 \tag{D.19}$$

We can now return to the problem of the planner. Using our results so far, we can write planner preferences as

$$\sum_{t=0}^{\infty}\beta^t U_t \approx -\nu_\ell(\bar{\ell})\bar{\ell}\sum_{t=0}^{\infty}\beta^t\left[\widehat{d}_t + \frac{1}{2}(\gamma+\phi)\,\widehat{y}_t^2 + \frac{\gamma}{2}\int\frac{\widehat{\omega}(\zeta^t)^2}{\bar{\omega}(\zeta^t)}d\Gamma(\zeta^t)\right]$$

Note that $\widehat{\pi}_t^2$ affects $\widehat{d}_t, \beta\widehat{d}_{t+1}, \cdots$ by $\frac{\theta\bar{\eta}}{2(1-\theta)}(1, \beta\theta, \cdots)$ so the discounted sum is $\frac{\theta\bar{\eta}}{2(1-\theta)(1-\beta\theta)}$. Using this we have

$$\begin{aligned}\sum_{t=0}^{\infty}\beta^t U_t &\approx -\nu_\ell(\bar{\ell})\bar{\ell}\sum_{t=0}^{\infty}\beta^t\left[\frac{\theta\bar{\eta}}{2(1-\theta)(1-\beta\theta)}\widehat{\pi}_t^2 + \frac{1}{2}(\gamma+\phi)\,\widehat{y}_t^2 + \frac{\gamma}{2}\int\frac{\widehat{\omega}(\zeta)^2}{\bar{\omega}(\zeta)}d\Gamma(\zeta)\right] \\ &= -\frac{\nu_\ell(\bar{\ell})\bar{\ell}\theta\bar{\eta}}{2(1-\theta)(1-\beta\theta)}\sum_{t=0}^{\infty}\beta^t\left[\widehat{\pi}_t^2 + \frac{\kappa}{\bar{\eta}}\widehat{y}_t^2 + \frac{\kappa\gamma}{(\gamma+\phi)\,\bar{\eta}}\int\frac{\widehat{\omega}(\zeta^t)^2}{\bar{\omega}(\zeta^t)}d\Gamma(\zeta^t)\right],\end{aligned} \tag{D.20}$$

where we have used the definition of $\kappa = (\phi+\gamma)(1-\theta)(1-\beta\theta)/\theta$. □

D.4 Auxiliary lemma for Proposition A.1

We here establish that, under Assumptions A.1 to A.3, changes in nominal interest rates do not affect the distribution of consumption shares. We proceed in two steps. First, we show that consumption dispersion is unaffected by changes in real interest rates and output that satisfy particular conditions. Second, we establish that changes in the monetary policy stance induce changes in real interest rates and output that satisfy precisely those conditions. Throughout, our arguments closely follow Werning (2015).

Lemma D.1. *Suppose that Assumptions A.1 to A.3 hold, and consider paths* $(\boldsymbol{r}, \boldsymbol{y}, \boldsymbol{m}, \boldsymbol{\tau}_x)$ *such that* $\widehat{\boldsymbol{m}} = \boldsymbol{\tau}_x = \mathbf{0}$ *and, for all* $t = 0, 1, 2, \ldots,$

$$y_t^{-1} = \tilde{\beta}(1 + r_{t+1}) y_{t+1}^{-1}, \tag{D.21}$$

where $\tilde{\beta} \equiv (1 + \bar{r})^{-1}$. *Then the distribution of consumption shares remains constant at its steady state distribution:* $\omega_t(\zeta) = \bar{\omega}(\zeta)$ *for all* t *and* ζ.

Proof. We will rewrite the household budget constraint using several substitutions. Real aggregate labor income is equal to aggregate output, i.e., $\ell_t w_t / p_t = y_t$. The government budget constraint is $\bar{r} y_t \bar{b} + \tau_{e,t} = \tau_y y_t$, where $\bar{b}$ is the constant level of government debt outstanding (measured as a number of bonds whose price will fluctuate), and we can use this to substitute out for $\tau_{e,t}$. Putting the pieces together, the household budget constraint (8) becomes

$$a_{it} + c_{it} = (1 + r_t) a_{it-1} + \Phi(\zeta_{it}, \bar{m}) y_t - \bar{r} y_t \bar{b} e_{it}, \tag{D.22}$$

where we have used $e_{it} = \Phi(\zeta_{it}, \bar{m})$. In this budget constraint, a_{it} is the value of savings measured in terms of the final good. Re-write this as $a_{it} = q_t b_{it}$ where b_{it} is the number of perpetuities purchased by the household, each of which trades at a price of q_t (denominated in final goods). Similarly $(1 + r_t) a_{it-1}$ is the value of asset position at the start of the period, which we can re-write as $(1 + r_t) a_{it-1} = \bar{r} y_t / \bar{y} b_{it-1} + q_t b_{it-1}$, where the first term is the "coupon" and the second term is the value of the previously held assets. Combining the asset-pricing relationship (13) and the aggregate Euler equation (D.21), we have

$$\frac{q_t}{y_t} = \tilde{\beta} \left(\frac{\bar{r}}{\bar{y}} + \frac{q_{t+1}}{y_{t+1}} \right).$$

Provided that steady state real rates are positive, we have that $\tilde{\beta} < 1$ and so we can solve forward to find $q_t = \tilde{\beta} \bar{r} (y_t / \bar{y}) / (1 - \tilde{\beta}) = y_t / \bar{y}$. After substituting for a_{it} and a_{it-1}, the household budget constraint becomes

$$\frac{y_t}{\bar{y}} b_{it} + c_{it} = (1+\bar{r}) \frac{y_t}{\bar{y}} b_{it-1} + \left(1 - \bar{r}\bar{b}\right) \Phi(\zeta_{it}, \bar{m}) y_t. \tag{D.23}$$

Letting z denote an arbitrary realization of ζ_{it}, we will now re-state the household consumption-savings problem in recursive form. We have

$$V_t(b,z) = \max_{b' \geq 0} \log \left[(1+\bar{r})b + \left(1 - \bar{r}\bar{b}\right) \Phi(z, \bar{m}) - b'\right] + \log\left(\frac{y_t}{\bar{y}}\right) + \beta \mathbb{E}\left[V_{t+1}(b', z')\right]$$

If V_{t+1} is of the form $V_{t+1}(b,z) = \tilde{V}(b,z) + B_{t+1}$ for some sequence B_{t+1}, then the Bellman equation above can be written as

$$V_t(b,z) = \max_{b' \geq 0} \log \left[(1+\bar{r})b + (1 - \bar{r}\bar{a}) \Phi(z, \bar{m}) - b'\right] + \beta \mathbb{E}\left[\tilde{V}(b', z')\right] + B_t, \tag{D.24}$$

where $B_t = \log(y_t/\bar{y}) + \beta B_{t+1}$. As there is no time-varying aggregate variable apart from B_t, V_t satisfies the same functional form as V_{t+1}. By induction, all previous value functions satisfy this form. Using the steady-state value function to start the induction (i.e., we start at the steady-state value function $\tilde{V}(b,z)$ and $B_t = 0$), we can conclude from (D.24) that the optimal decision rule for b' as a function of (b,z) will be constant across time. This constant decision rule and a stable process for the evolution of z' implies the distribution of (b', z') is unaffected. It follows from (D.23) that the optimal consumption decision rule will scale with $y_t = c_t$. This scaling implies consumption shares are constant and equal to their steady state values.

As a final step, it remains to relate this recursive formulation of the household decision problem to the histories of idiosyncratic events. To this end, note that we can write the consumption share as a function of the state variables associated with that history:

$$\omega_t(\zeta) = \frac{c(b(\zeta), z(\zeta))}{\bar{c}}, \tag{D.25}$$

where $c(b,z)$ is the steady state consumption function, $b(\zeta)$ is the steady state bond holdings of a household with history ζ and $z(\zeta)$ is the most recent event in the history ζ. (D.25) holds for any paths $(\boldsymbol{r}, \boldsymbol{y}, \boldsymbol{m}, \boldsymbol{\tau}_x)$ such that (D.21) holds and $\widehat{\boldsymbol{m}} = \boldsymbol{\tau}_x = \mathbf{0}$. □

D.5 Proof of Proposition A.1

It remains to show that changes in nominal interest rates induce paths of $\boldsymbol{r}$ and $\boldsymbol{y}$ that satisfy (D.21) (since by linearity we already have $\widehat{\boldsymbol{m}} = \boldsymbol{\tau}_x = \mathbf{0}$). But this follows directly from Werning (2015), as our model economy with Assumptions A.1 to A.3 satisfies the conditions of his result (i.e., acyclical risk and acyclical liquidity). We furthermore note that our special case is isomorphic to the incomplete markets model that appears in Section IIIB of Farhi & Werning (2019).[33] We refer the reader to Werning (2015) for the formal proof. □

D.6 Proof of Corollary A.1

By (4) we can write the optimal monetary policy rule as

$$\Theta'_{\pi,i} W \widehat{\boldsymbol{\pi}} + \frac{\kappa}{\bar{\eta}} \Theta'_{y,i} W \widehat{\boldsymbol{y}} = \mathbf{0} \tag{D.26}$$

It follows from (D.1) that

$$\Pi_\pi \Theta_{\pi,i} = \Pi_y \Theta_{y,i} \tag{D.27}$$

and so we can re-write (D.26) as

$$\widehat{\boldsymbol{\pi}} + \frac{\kappa}{\bar{\eta}} \frac{1}{\kappa} \begin{pmatrix} 1 & 0 & 0 & \dots \\ -1 & 1 & 0 & \dots \\ 0 & -1 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \widehat{\boldsymbol{y}} = \mathbf{0} \tag{D.28}$$

giving (A.5). □

D.7 Proof of Corollary A.2

We have already shown that the optimal monetary rule is given as (A.5). Since (D.1) also implies that

$$\Pi_\pi \Theta_{\pi,\tau_x} = \Pi_y \Theta_{y,\tau_x} \tag{D.29}$$

[33] The model in Farhi & Werning (2019) specifies a particular AR(1) process for idiosyncratic income risk for the sake of computing numerical solutions. We leave the process more general. The important aspect is that risk is not affected by monetary policy (see Werning, 2015).

we can use the same steps as in the proof of Corollary A.1 to rewrite the first two terms in (34) as

$$\Theta'_{\pi,\tau_x} W \widehat{\boldsymbol{\pi}} + \frac{\kappa}{\bar{\eta}} \Theta'_{y,\tau_x} W \widehat{\boldsymbol{y}} = \widehat{\boldsymbol{\pi}} + \frac{1}{\bar{\eta}} \begin{pmatrix} 1 & 0 & 0 & \dots \\ -1 & 1 & 0 & \dots \\ 0 & -1 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \widehat{\boldsymbol{y}}.$$

Imposing (A.5) sets these terms to zero, so Corollary A.2 follows. $\square$