Optimal Policy Rules in HANK[†]

Alisdair McKay
FRB Minneapolis

Christian K. Wolf
MIT & NBER

Preliminary and incomplete – comments welcome!

January 6, 2022

Abstract: We propose a new method to characterize optimal policy rules in business-cycle models with nominal rigidities and heterogeneous households. Our approach applies to simple loss functions (e.g., a dual mandate) and to the full Ramsey problem alike, comes at little additional computational cost, and yields optimal policy rules in the form of forecast targeting criteria. Applied to optimal stabilization policy in a quantitative business-cycle model, this method yields three main insights. First, household heterogeneity has no effect on the target criterion of a dual mandate policymaker. Second, the Ramsey interest rate target criterion materially differs from the dual mandate rule if and only if interest rate policy has large effects on consumption inequality. Third, in response to adverse shocks concentrated at the bottom of the income distribution, stimulus checks generally achieve a much smaller loss than conventional interest rate policy.

[†]Email: alisdair.mckay@mpls.frb.org and ckwolf@mit.edu. 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

Modern monetary policy strategies are typically expressed in terms of forecast target criteria—an implicit rule prescribing that policy instruments are set so that economic outcomes are expected to be in line with policymaker targets. For example, an inflation-targeting central bank sets interest rates so that inflation is expected to return to its target over time. Such target criteria are popular in policymaking practice because they are simple to communicate (Svensson, 1997), yet also theoretically appealing: in standard representative-agent business-cycle models, well-chosen target criteria robustly implement social optima, independently of the shocks hitting the economy (Giannoni & Woodford, 2002).

In this paper we connect the literature on forecast targeting criteria with recently growing interest in the positive and normative effects of inequality on policy design. To do so, we develop a novel computational approach that yields optimal policy rules in heterogeneous-agent models with nominal rigidities ("HANK"), expressed in the form of forecast target criteria. We then use our approach to characterize optimal policy rules in HANK for two leading cases of the policymaker objective—a simple dual mandate and a full Ramsey planner loss function—and two popular stabilization policy tools—short-term nominal interest rates and stimulus checks. The main payoff of our analysis is to clarify the conditions under which inequality affects optimal target criteria, and if so, how.

The key conceptual innovation underlying our approach is to consider optimal policy problems that can be written as a linear-quadratic optimization problem in *sequence-space*, requiring as inputs only: (i) linearized aggregate equilibrium conditions; and (ii) a quadratic loss function, either mandated from outside or derived as an approximation to a social welfare function. We show that the first-order conditions of this problem immediately give us optimal rules in the form of forecast targeting criteria.[1] We furthermore establish that these target criteria depend on the structure of the underlying model *only* through the dynamic causal effects of policy instruments on policymaker targets—a robust set of "sufficient statistics" characterizing optimal policy rules. Our approach is computationally cheap precisely because those causal effects can be obtained straightforwardly through well-established methods that solve for equilibria *given* policy rules (e.g. Auclert et al., 2021).

We showcase our approach by solving for optimal policy rules in a canonical business-cycle model with nominal rigidities and household heterogeneity. Households face idiosyncratic

[1] 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.

income risk and self-insure by borrowing and saving in a (long-duration) bond, 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. While our derived target criteria will apply independently of the underlying shock processes, we pay particular attention to policy responses to two particular shocks: a wage cost-push "supply" shock and an earnings inequality "demand" shock.

We show that optimal policy problems in this environment can be cast in the form of our general linear-quadratic set-up. Formally, our objective is to evaluate policymaker welfare to second-order through a first-order (linearized) approximation of aggregate dynamics. For the second-order approximation of policymaker welfare we consider two different options.

1. A (quadratic) loss function may simply be legislated to the policymaker, e.g. in the form of a dual mandate. We regard this case as particularly relevant for monetary policy design because of its similarity to the legislated mandates of central banks.

2. Policymaker preferences may be given as a (weighted) average of household utilities. For a linearized equilibrium representation to allow evaluation of such preferences to second order, we require the model's deterministic steady state to correspond to a policymaker optimum (Giannoni & Woodford, 2002). In principle this can be ensured either by setting the long-run tax-and-transfer system optimally given policymaker preference weights, or by setting the weights to rationalize a given tax-and-transfer system. We follow the second route and then, given efficiency, compute our second-order approximation of policymaker preferences. We find that it consists of three terms: the usual output and inflation terms, and a novel "inequality" term, reflecting fluctuations in the distribution of consumption.

Linearizing the model's non-policy block, we obtain three constraints: a generalized "IS"-curve as the demand block; a standard Phillips curve as the supply block; and the evolution of the consumption distribution. The computational simplicity of our method follows precisely because these three constraints can be derived either immediately or with little additional work from standard sequence-space solution output (Auclert et al., 2021). Thus we now have all the pieces in place to leverage the results from our general linear-quadratic problem.

We first study optimal policy for an externally mandated "dual mandate"-type loss function. Our main finding here is that the optimal interest rate target criterion is *exactly the same* as in a 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 always a slack constraint: the policymaker picks a desired output-inflation allocation subject to the Phillips curve, and then simply sets nominal interest rates to whatever level necessary to generate aggregate demand consistent with the desired allocation. Household heterogeneity thus matters for the path of the policy instrument, but not for the forecast target criterion that this instrument implements, and so in particular not for the equilibrium paths of output and inflation. Leveraging the results in Wolf (2021), we furthermore show that the optimal policy problem for *joint* interest rate and transfer setting features an indeterminacy: given the lack of Ricardian equivalence, the policymaker can engineer the required path of demand either through nominal interest rates or through transfer payments, with the split between the two left undetermined.

We then turn to the general Ramsey problem. With inequality now a policymaker target, our general formulas for optimal targeting criteria imply that inequality-related terms will appear if and only if the policy instrument affects consumption inequality. To illustrate this point, we follow Werning (2015) and consider a benchmark case of our economy in which interest rate policy does not affect cross-sectional consumption dispersion. In that case, even though the planner intrinsically cares about inequality, all inequality-related terms drop out of the forecast targeting criterion for optimal monetary policy, so the criterion looks exactly as in the simple dual mandate case. For example, after a contractionary demand impulse related to an increase in earnings inequality, the policymaker still cuts rates to stabilize total consumption, with no ability to offset the shock's distributional effects. This contrasts with transfer payments, which invariably do affect consumption inequality. This benchmark Werning (2015) case thus yields a clean separation between the two policy tools: transfers are used to manage inequality objectives, while monetary policy is set residually to bring total demand in line with target. In future versions of this paper, we plan to quantify these effects on optimal policy rules in a version of the model estimated to be consistent with empirical evidence on the distributional incidence of monetary and transfer policies.

Literature. Our work relates and contributes to several strands of literature.

First, our methodology for optimal policy analysis extends the linear-quadratic approach popular in representative-agent settings (Woodford, 2011) to models with rich micro heterogeneity. This is useful for two reasons: first, we import all of the advantages of forecast targeting rules to optimal policy analysis in the HANK environment; and second, our solution is computationally straightforward, leveraging recent advances of Auclert et al. (2021).[2]

[2]Sequence-space linear-quadratic policy problems have been used in prior work to derive "optimal policy

Second, our results speak to the question of how and why inequality should inform optimal policy design. Existing studies of optimal policy in HANK models are either largely numerical (Bhandari et al., 2021; Le Grand et al., 2021) or study environments with limited heterogeneity (Acharya et al., 2020). While our method allows for straightforward computation of optimal policies, it also yields useful analytical insights that apply even in an environment with rich heterogeneity. In particular we find that, for conventional dual mandate policymakers, inequality will in general affect the required paths of policy instruments, but not the optimal evolution of output and inflation. For forecast targeting central banks, this means that there is no need to change operational procedures: instruments can continue to be set in whatever way necessary to align projections of macroeconomic aggregates with the mandated targets. For the full Ramsey problem, we show that a third term enters the policymaker loss, and give conditions under which that term will also shape optimal forecast targeting criteria. Finally, our results on optimal joint fiscal-monetary policy in the presence of this inequality term extend previous 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.

Third, we connect our derived optimal forecast target criteria to empirical measurement. As we emphasize in McKay & Wolf (2021), empirical evidence on the dynamic causal effects of policy shocks at the same time pins down individual entries of the optimal policy rules presented here.[3] Our results on the general Ramsey problem thus suggest that evidence on the distributional implications of monetary policy *shocks* (Coibion et al., 2017; Slacalek et al., 2020; Andersen et al., 2021) will be useful to further pin down optimal Ramsey monetary

projections" (Svensson, 2005; De Groot et al., 2021; Hebden & Winker, 2021). We make two contributions relative to those studies. First, we show how to fit optimal policy problems in HANK into this environment, including in particular the derivation of a quadratic loss function for the Ramsey problem. Second, we write the solution to the sequence-space policy problem in terms of target criteria that are expressed as functions of (empirically measurable) shock impulse response functions of target variables.

[3]In McKay & Wolf (2021), we show that, *given a loss function*, empirical time series evidence is in principle sufficient to fully characterize optimal policy rules. In this paper we instead *derive* loss functions in a model and show how to *compute* optimal policy rules. The insights from our other paper are then useful to inform the model calibration.

policy *rules*. In future versions of this paper we plan to precisely use such empirical evidence to discipline a quantitative version of our structural model.

Outline. The remainder of the paper proceeds as follows. In Section 2 we present a general linear-quadratic sequence-space problem and its solution. Section 3 then outlines our particular structural model, and Section 4 shows that optimal policy problems in that model can be cast in our convenient linear-quadratic form. Results are presented in Section 5. We conclude in Section 6 and relegate supplementary results to several appendices.

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 results in the broader context of the literature, discussing in particular the relationship between our deterministic (perfect-foresight) set-up and the stochastic linear-quadratic problem studied in much prior work (for example Benigno & Woodford, 2012).

2.1 Problem & solution

We begin by describing the linear-quadratic optimal control problem of the policymaker.

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

$$\mathcal{L} \;\equiv\; \frac{1}{2}\sum_{i=1}^{I} \lambda_i \boldsymbol{x}_i' W \boldsymbol{x}_i \;=\; \frac{1}{2}\boldsymbol{x}' \left(\Lambda \otimes W\right) \boldsymbol{x}, \tag{1}$$

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 $\Lambda \equiv \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_I)$. Finally $W = \text{diag}(1, \beta, \beta^2, \dots)$ 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

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

where $\boldsymbol{z} \equiv (\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_J)'$ stacks time paths for the J policy instruments available to the policymaker, and $\boldsymbol{\varepsilon} \equiv (\boldsymbol{\varepsilon}_1, \boldsymbol{\varepsilon}_2, \ldots, \boldsymbol{\varepsilon}_Q)'$ similarly stacks the paths for Q exogenous "shocks."[4] $\{\mathcal{H}_x, \mathcal{H}_z, \mathcal{H}_\varepsilon\}$ 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 linear-quadratic optimal policy problem is now 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

[4] Note that the exogenous shocks can include the effects of initial conditions for the state variables (in addition to the more usual form of shocks as exogenous driving forces).

[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 B.1 for further details.

satisfy the "policy criterion"

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

(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 (2011), this rule is an example of a so-called implicit "target policy criterion": the policymaker sets the available instruments to align projections (i.e., future paths) of macro aggregates as well as possible with its targets, given what is achievable through the available instruments.

We emphasize two important features of such rules. First, they are derived without reference to and so apply independently of the actual 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

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

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 $\Theta_{x,\varepsilon} \cdot \boldsymbol{\varepsilon}$. 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 implies 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 explore this connection further in our quantitative analysis in Section 5.3.

Outlook. In the remainder of this paper we will first show that optimal policy problems in models with rich microeconomic heterogeneity can be represented in the form of our

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

linear-quadratic control problem, and then use (4) and (5) to characterize optimal policy rules. Sections 3 and 4 begin with the first part of that argument.

3 Model

Our model environment is a standard Heterogeneous-Agent New Keynesian ("HANK") economy, with two somewhat special features. First, our model features sticky wages. While the early HANK literature focussed on sticky price models (e.g., 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 in our quantitative analysis.

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}

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

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 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 analysis in Section 5, 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 thus refer to it as the "exogenous" component of transfers, hence the subscript x. 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 household 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 the derivative matrices $\mathcal{C}_{\bullet}$ are evaluated at 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 k, ε_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 a 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 that everyone supplies an equal amount of hours worked, and with our household preference specification in (7) additively separable, 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 (NKPC):

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

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 $\psi \equiv -\frac{\kappa}{(\phi+\gamma)(\epsilon-1)}$. In our derivation of (11), 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 4.1, 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, 2011).

Finally, aggregate production is equal to

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

where $\ell_t \equiv \int_0^1 \int_0^1 e_{it}\ell_{ikt} di dk$ and the term d_t captures efficiency losses related to wage dispersion across unions.

3.3 Asset structure

There are two different assets in the economy: a short-term, risk-free nominal bond 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 provide the same expected returns along equilibrium transition paths (except at $t = 0$), 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—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 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.

Let r_t denote the real return on the second asset. Bond coupon payments are 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.

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 $\tau_{x,t}$ as the independent policy instruments of the government, used for business-cycle stabilization policy. Transfers are deficit-financed, with government debt evolving according to the rule

$$b_t = (1 - \rho_b)\bar{b} + \rho_b b_{t-1} + \tau_{x,t}. \tag{16}$$

The endogenous component of transfers $\tau_{e,t}$ then adjusts residually to balance the government budget at all t.

3.5 Equilibrium

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

Definition 1. *Given paths of exogenous shocks* $\{m_t, \varepsilon_t\}_{t=0}^{\infty}$*, 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* $\{\pi_t, y_t, \varepsilon_t\}_{t=0}^{\infty}$ *are consistent with the aggregate NKPC* (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* $a_t = q_t b_t$*.*

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

Discussion. How does our model differ from the canonical representative agent New Keynesian models (Galí, 2015; Woodford, 2011)? *Positively*, the main change is that a simple aggregate Euler equation,

$$\widehat{c}_t = -\frac{1}{\gamma}\widehat{r}_{t+1} + \widehat{c}_{t+1}, \tag{17}$$

is now replaced by the more general aggregate consumption function (10). Inequality thus affects the aggregate dynamics of our economy in response to shocks and policy actions only

[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). Identical labor rationing across households then implies $\ell_{it} = \ell_t$.

through the demand side, with supply—and so in particular the Phillips curve (11)—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. We will discuss these changes in Section 4.1.

4 The optimal policy problem

This section presents the optimal policy problem, casting it in a form consistent with our linear-quadratic sequence-space set-up in Section 2. Section 4.1 begins with the (quadratic) loss function, Section 4.2 reduces the model economy of Section 3 to a small number of linear constraints, and Section 4.3 summarizes the resulting optimal policy problem.

4.1 Loss function

We consider two loss functions: an externally mandated dual mandate loss function and a quadratic approximation to a general social welfare function for the model in Section 3.

DUAL MANDATE. Our first loss function is not derived from any primitives of the model in Section 3, but instead assumed outright, taking the simple form

$$\mathcal{L}^{DM} \equiv \sum_{t=0}^{\infty} \beta^t \left[\lambda_\pi \widehat{\pi}_t^2 + \lambda_y \widehat{y}_t^2\right]. \tag{18}$$

(18) 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. While ad-hoc, we regard such a loss function as interesting because of its clear practical relevance, with central banks across the world receiving legislatively mandated policy targets similar to (18).

[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.

Ramsey planner. For our second loss function, we consider a conventional Ramsey planner that aggregates utilities of the households populating the economy. 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 $\omega_t(\zeta_i^t)$ is their share of aggregate consumption.[11] 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

$$\mathcal{V}^{HA} = \sum_{t=0}^{\infty} \beta^t \int \varphi(\zeta) \left[\frac{(\omega_t(\zeta)c_t)^{1-\gamma} - 1}{1-\gamma} - \nu(\ell_t) \right] d\Gamma(\zeta), \tag{19}$$

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, 2011), our objective is to evaluate the social welfare function (19) to second order. To this end, a first-order approximation to aggregate equilibrium dynamics—as embedded in our linearized transition path equilibrium concept in Definition 1—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-state distribution of consumption given a tax-and-transfer system is in fact optimal. In this paper we adopt the second approach, following much of the inverse optimal taxation literature (e.g. Heathcote & Tsujiyama, 2021). Our preference for this solution 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 C.1 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 (19) around the efficient steady state then gives the following characterization of the policymaker loss function.

Proposition 1. *To second order, the social welfare function* (19) *is proportional to* $-\mathcal{L}^{HA}$*,*

[11] 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.

given as

$$\mathcal{L}^{HA} \equiv \sum_{t=0}^{\infty} \beta^t \left[\widehat{\pi}_t^2 + \frac{\kappa}{\bar{\varepsilon}} \widehat{y}_t^2 + \frac{\kappa\gamma}{(\gamma+\phi)\bar{\varepsilon}} \int \frac{\widehat{\omega}_t(\zeta)^2}{\bar{\omega}(\zeta)} d\Gamma(\zeta) \right], \tag{20}$$

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, 2011). 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 (20) fit into the linear quadratic framework in Section 2? Moving to a sequence-space formulation, we can write the loss as

$$\mathcal{L}^{HA} = \lambda_\pi \widehat{\boldsymbol{\pi}}' W \widehat{\boldsymbol{\pi}} + \lambda_y \widehat{\boldsymbol{y}}' W \widehat{\boldsymbol{y}} + \int \lambda_{\omega(\zeta)} \widehat{\boldsymbol{\omega}}(\zeta)' W \widehat{\boldsymbol{\omega}}(\zeta) d\Gamma(\zeta), \tag{21}$$

where $\lambda_\pi = 1$, $\lambda_y = \frac{\kappa}{\bar{\varepsilon}}$ and $\lambda_{\omega(\zeta)} \equiv \frac{\kappa\gamma}{(\gamma+\phi)\bar{\varepsilon}\bar{\omega}(\zeta)}$. 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 (21) in Section 4.3.[12]

4.2 Constraints

Next, we summarize the rich model of Section 3 through a small number of linear relations—the eventual constraints of the optimal policy problem. We begin with a simple result that provides a more compact characterization of equilibrium dynamics.

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

$$\widehat{\boldsymbol{\pi}} = \kappa \widehat{\boldsymbol{y}} + \beta \widehat{\boldsymbol{\pi}}_{+1} + \psi \widehat{\boldsymbol{\varepsilon}}, \tag{22}$$

$$\widehat{\boldsymbol{y}} = \tilde{\mathcal{C}}_y \widehat{\boldsymbol{y}} + \tilde{\mathcal{C}}_\pi \widehat{\boldsymbol{\pi}} + \tilde{\mathcal{C}}_i \widehat{\boldsymbol{i}} + \tilde{\mathcal{C}}_x \widehat{\boldsymbol{\tau}}_x + \mathcal{C}_m \widehat{\boldsymbol{m}}, \tag{23}$$

[12] Note that, technically, (21) 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 discussed in Section 4.3 will consider an equivalent formulation of the problem with finitely many policy targets.

where the linear maps $\{\tilde{\mathcal{C}}_y, \tilde{\mathcal{C}}_\pi, \tilde{\mathcal{C}}_i, \tilde{\mathcal{C}}_x\}$ *are defined in Appendix C.2, and* $\widehat{\boldsymbol{\pi}}_{+1} = (\pi_1, \pi_2, \dots)$.

Lemma 1 reduces the complexity of the equilibrium in Section 3 to two equations: the Phillips curve (22) (which is simply a stacked perfect-foresight version of the original relation (11)); and the IS curve (23), 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, $\tau_{e,t}$. Equation (23) is the natural analogue of the IS-curve constraint considered for optimal analysis in standard representative-agent models (see Woodford, 2011).

Equations (22) - (23) fully characterize the evolution of output and inflation given exogenous non-policy shocks and policy choices. It remains to describe the evolution of the inequality term (20) as a function of shocks, policies, and the equilibrium paths of aggregate output and inflation. In Appendix B.2, we establish that, to first order, we can write the consumption share for a household with specific history ζ as

$$\widehat{\boldsymbol{\omega}}(\zeta) = \Omega_{\omega(\zeta),y}\widehat{\boldsymbol{y}} + \Omega_{\omega(\zeta),\pi}\widehat{\boldsymbol{\pi}} + \Omega_{\omega(\zeta),i}\widehat{\boldsymbol{i}} + \Omega_{\omega(\zeta),x}\widehat{\boldsymbol{\tau}}_x + \Omega_{\omega(\zeta),m}\widehat{\boldsymbol{m}}, \qquad \forall\zeta \tag{24}$$

where the maps $\Omega_\bullet$ give the derivatives of consumption shares with respect to aggregate variables. The intuition for (24) 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 shocks, policies, and output and inflation paths.

4.3 Summary & computational details

We now have all the ingredients required to state our two optimal policy problems in the notation of the general linear-quadratic sequence-space problem of Section 2.

Dual mandate. The dual mandate policymaker chooses paths of nominal interest rates $\boldsymbol{i}$ and the exogenous component of transfers $\boldsymbol{\tau}_x$ to minimize the mandated loss function $\mathcal{L}^{DM}$ in (18) subject to the constraints (22) and (23): the Phillips curve and the generalized IS curve. The resulting 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 (17) to the rich demand relation (23).

Note that so far the constraints of this policy problem take the form of our general linear 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).[13] For future reference, we write these impulse response space constraints as

$$\widehat{\boldsymbol{\pi}} = \Theta_{\pi,i}\widehat{\boldsymbol{i}} + \Theta_{\pi,x}\widehat{\boldsymbol{\tau}}_x + \Theta_{\pi,\varepsilon}\widehat{\boldsymbol{\varepsilon}} + \Theta_{\pi,m}\widehat{\boldsymbol{m}}, \tag{25}$$

$$\widehat{\boldsymbol{y}} = \Theta_{y,i}\widehat{\boldsymbol{i}} + \Theta_{y,x}\widehat{\boldsymbol{\tau}}_x + \Theta_{y,\varepsilon}\widehat{\boldsymbol{\varepsilon}} + \Theta_{y,m}\widehat{\boldsymbol{m}}. \tag{26}$$

Key to the computational simplicity of our approach to optimal policy analysis 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 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 obtaining an *optimal* policy rule.

Ramsey planner. The full Ramsey planner chooses paths of the same two 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, (22) and (23), as well as the evolution of the inequality term, (24). This problem now departs from previous optimal policy analyses in the representative-agent context through *both* the demand constraint and the introduction of a new target related to cross-sectional consumption dispersion. Re-expressed in equivalent impulse response space (again following McKay & Wolf (2021)), the evolution of consumption shares in the constraint (24) can alternatively be written as

$$\widehat{\boldsymbol{\omega}}(\zeta) = \Theta_{\omega(\zeta),i}\widehat{\boldsymbol{i}} + \Theta_{\omega(\zeta),x}\widehat{\boldsymbol{\tau}}_x + \Theta_{\omega(\zeta),\varepsilon}\widehat{\boldsymbol{\varepsilon}} + \Theta_{\omega(\zeta),m}\widehat{\boldsymbol{m}}. \qquad \forall\zeta \tag{27}$$

[13] 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 $\Theta_{\pi,i} \equiv \tilde{\Theta}_{\pi,i}\tilde{\Theta}_{i,i}^{-1}$. The arguments in McKay & Wolf (2021) reveal that this re-writing is without loss of generality.

To computationally evaluate the more complicated loss function (21) and the associated constraints, we show in Appendix B.2 that the Ramsey loss can be re-written as

$$\mathcal{L}^{HA} = \widehat{\boldsymbol{\pi}}' W \widehat{\boldsymbol{\pi}} + \frac{\kappa}{\bar{\varepsilon}} \widehat{\boldsymbol{y}}' W \widehat{\boldsymbol{y}} + \frac{\kappa\gamma}{(\gamma+\phi)\bar{\varepsilon}} \widehat{\boldsymbol{x}}' Q \widehat{\boldsymbol{x}} \tag{28}$$

where $\boldsymbol{x} = (\boldsymbol{y}, \boldsymbol{r}, \boldsymbol{\tau}_x, \boldsymbol{\tau}_e, \boldsymbol{m})$ and Q is a linear map, defined in Appendix B.2.[14] The alternative formulation in (28) reflects the simple intuition that it is always possible to re-write the evaluation of 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 (28), 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 B.2 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 Results

This section presents our results on optimal policy rules. As we have seen in Sections 3 and 4, household heterogeneity in principle has both positive and normative implications for optimal stabilization policy: positively, it changes the propagation of shocks and policies to macro aggregates through its effects on the demand side of the economy; normatively, it changes policymaker targets.

Our analysis in this section disentangles these forces. We begin in Section 5.1 with the optimal policy problem of a dual mandate policymaker, with heterogeneity mattering only through its effects on policy and shock propagation. Sections 5.2 and 5.3 then consider the full Ramsey problem, first in an instructive special case and then in a quantitative calibration. Throughout, we derive optimal rules for both interest rate and stimulus check policies.

[14]The linear map Q in (28) is in general not diagonal. Thus Appendix B.2 also extends the results from Section 2 to the case with such interaction terms.

5.1 Optimal dual mandate policy

We begin with the optimal monetary policy problem of a dual mandate central banker. Our first main result is that, under mild regularity conditions on the linear map $\tilde{\mathcal{C}}_i$—the mapping from nominal interest rate policy to net excess consumption demand in the generalized "IS" curve (23)—, the optimal monetary policy rule in the form of a forecast targeting criterion is completely unaffected by household heterogeneity.

Proposition 2. *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*

$$\widehat{\boldsymbol{c}} \in \textit{image}(\tilde{\mathcal{C}}_i), \tag{29}$$

then the optimal monetary policy rule for a dual mandate policymaker with loss function (18) *can be written as the forecast target criterion*

$$\lambda_\pi \widehat{\pi}_t + \frac{\lambda_y}{\kappa}\left(\widehat{y}_t - \widehat{y}_{t-1}\right) = 0, \quad \forall t = 0, 1, \ldots \tag{30}$$

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

$$\lambda_\pi \cdot \Theta'_{\pi,i} \cdot \boldsymbol{\pi} + \lambda_y \cdot \Theta'_{y,i} \cdot \boldsymbol{y} = \mathbf{0}. \tag{31}$$

The proof of Proposition 2 leverages the structure of our particular model to turn the general expression (31) into the simple rule (30). Importantly, the rule (30) is *exactly* the same as in conventional representative-agent optimal policy analyses (e.g. as in Woodford, 2011). 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 according to the Phillips curve, and then sets interest rates residually to deliver the required time path of real demand. Our technical condition in (29) 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 2 reveals that this is sufficient to ensure that any desired output-inflation pair consistent with the Phillips curve (22) 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 (29) 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 general feature of HANK-type environments.

Implications for monetary policy practice. The upshot of Proposition 2 is that, independently of the shock processes hitting the economy, under optimal dual mandate policy, the equilibrium paths of output and inflation will be unaffected by household heterogeneity and so equal to those in a representative-agent economy. The only effect of heterogeneity is to change the instrument paths—the current and future values of nominal interest rates—required to attain those desired output and inflation paths. The positive implications of household heterogeneity thus have a very 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 already done today.[15]

Optimal stimulus checks. Proposition 2 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

$$\widehat{\boldsymbol{c}} \in \text{image}(\tilde{\mathcal{C}}_\tau) \tag{32}$$

for all sequences $\widehat{\boldsymbol{c}}$ with zero net present value, then stimulus check policies can also implement the target criterion (30), just like conventional monetary policy. This alternative implementability condition (32) 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

[15] 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."

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 interest rates. The next two subsections consider an alternative resolution to this indeterminacy: a richer loss function.

5.2 Ramsey problem with Werning (2015) aggregation

We begin by studying the full Ramsey problem in an instructive special case of our economy: the analytical "as-if aggregation" case of Werning (2015). While inequality now enters the policymaker loss function, nominal interest rates will turn out to be powerless to actually *affect* inequality. Consistent with the results in Section 2, we will see that this again implies an irrelevance of heterogeneity for optimal monetary policy *rules*, with the forecast target criterion still taking the simple form (30).

Environment. 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 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 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 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 (17). 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.

Results. Following steps similar to those in Werning (2015), we can prove the following useful building block result.

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

$$\Theta_{\omega(\zeta),i} = \mathbf{0} \qquad \forall \zeta. \tag{33}$$

In words, nominal interest rates have no effect on the consumption *distribution*, at any horizon. Combining Proposition 3 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 1 summarizes this conclusion.

Corollary 1. *Under Assumptions 1 to 3 and the conditions of Proposition 2, the optimal monetary policy rule for a Ramsey policymaker with loss function* (20) *can be written as the forecast target criterion*

$$\widehat{\pi}_t + \frac{1}{\bar{\varepsilon}} \left(\widehat{y}_t - \widehat{y}_{t-1}\right) = 0, \quad \forall t = 0, 1, \ldots \tag{34}$$

(34) is the same rule as (30), with the only difference that the general weights (λ_π, λ_y) from before are now replaced with the weights implied by the policymaker's loss function (20). Corollary 1 reveals that *caring* about inequality is not enough to affect optimal policy rules—the policy instrument must also be able to materially *affect* inequality. 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.[16] The discussion in Section 5.3 will address the quantitative relevance of this benchmark.

Allowing for transfers. Before proceeding to the quantitative analysis, we conclude this section with a brief discussion of optimal *joint* transfer and interest rate policy. Unlike monetary policy, transfers even in this particular environment do affect the inequality term,

[16] Note that the equivalence to representative-agent environments is in this case even stronger than in Section 5.1: not only is the target criterion the same, it is also implemented by the exact same time path of nominal interest rates because the demand block of the economy continues to aggregate to a conventional Euler equation—Werning's original result.

and so inequality enters the transfer policy target criterion.[17] By the discussion in Section 2.1, this optimal target criterion for stimulus checks takes the form

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

The two policy criteria (34) - (35) fully characterize optimal joint monetary-fiscal policy. As we show in Corollary 2, the monetary policy target criterion in (34) sets the first two terms in (35) to zero, so it follows that we can equivalently write the transfer target criterion as

$$\int \lambda_{\omega(\zeta)}\Theta'_{\omega(\zeta),\tau_x} W\widehat{\boldsymbol{\omega}}(\zeta) d\Gamma(\zeta) = \mathbf{0}. \tag{36}$$

We thus achieve a strict separation of monetary and fiscal instruments.

Corollary 2. *Under the optimal joint fiscal-monetary policy, transfers are set following the target criterion* (36)*, minimizing the inequality term in the loss function* (20)*. Monetary policy is set residually to enforce the target criterion* (34)*, thus attaining the same paths of aggregate inflation and output as under optimal monetary policy alone.*

Corollary 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.

Figure 1 provides a brief numerical illustration for a calibrated version of our model.[18] The figure shows impulse responses of nominal interest rates (i_t), the exogenous component of transfers ($\tau_{x,t}$), and the quadratic inequality term ($\int \frac{\widehat{\omega}_t(\zeta)^2}{\bar{\omega}(\zeta)} d\Gamma(\zeta)$) to a distributional shock $\boldsymbol{m}$ under both (i) optimal monetary policy (grey) and (ii) optimal joint monetary-fiscal policy (blue). We consider a distributional shock that temporarily *increases* earnings inequality, thus depressing consumer spending. We normalize the size of the shock so that, under the optimal Ramsey monetary policy, nominal rates are cut by one percentage point on impact. Since this distributional shock $\boldsymbol{m}$ only enters the IS curve (23), it follows immediately from

[17]Note that the outstanding amount of government bonds follows (16), and so is allowed to fluctuate in response to transfer stimulus payments.

[18]For details of the model parameterization for this numerical illustration see Appendix A.2.

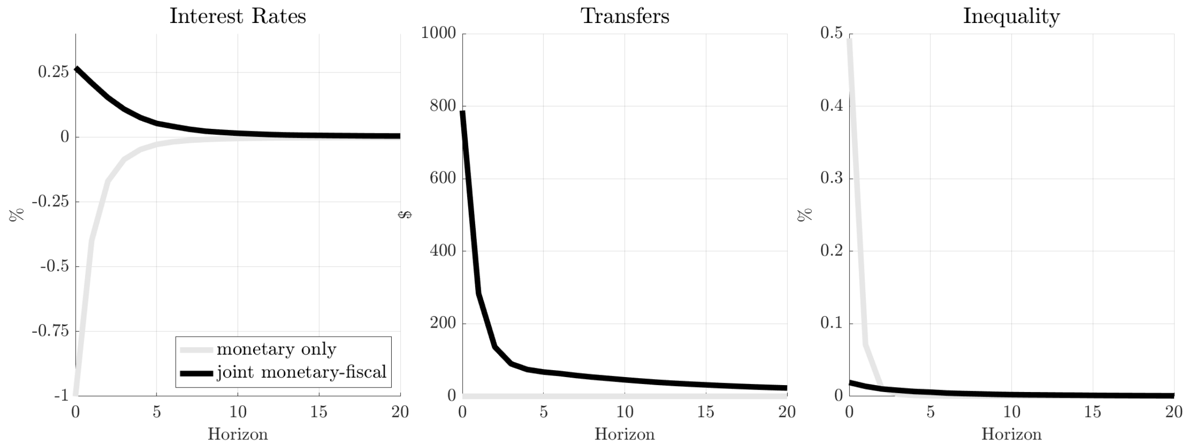

Figure 1: Impulse responses to an inequality shock under (i) optimal monetary (grey) and (ii) optimal joint monetary-fiscal policy (blue) in a calibrated version of our Werning model.

Corollary 1 and Corollary 2 that, for both sets of optimal policies, output and inflation are stabilized perfectly ("divine coincidence"). Monetary policy alone stabilizes demand through an aggressive interest rate cut, displayed in the left panel. Since this rate cut stimulates the consumption of all households equally, the policymaker loss related to the inequality term remains large, as shown in the right panel. Under the joint optimal policy, the policymaker instead sends out stimulus checks (middle panel), thus stabilizing consumption of the poor and bringing the inequality term closer to target (right panel). Combining the optimal target criterion (36) with our general results for optimal policy instrument paths derived in Section 2.1, we find this optimal transfer payment path as

$$\boldsymbol{\tau}_x = \left(\int \lambda_{\omega(\zeta)} \Theta'_{\omega(\zeta),\tau_x} W \Theta_{\omega(\zeta),\tau_x} d\Gamma(\zeta) \right)^{-1} \times \left(\int \lambda_{\omega(\zeta)} \Theta'_{\omega(\zeta),\tau_x} W \Theta_{\omega(\zeta),m} d\Gamma(\zeta) \cdot \boldsymbol{m} \right) \tag{37}$$

Transfers are thus set to undo the distributional effects of the inequality shock $\boldsymbol{m}$ as well as possible, in the usual weighted least-squares sense.[19] Monetary policy is then set residually to attain perfect output and inflation stabilization. Since the optimal transfer alone already provides a sizable temporary stimulus to demand, the monetary authority now attains the divine coincidence outcome through a less expansionary monetary policy (left panel); in fact, in our baseline calibration, nominal rates are in fact instead *increased.*

[19] Note that it also follows from the logic of the Werning aggregation result that, in our model, the supply shock $\boldsymbol{\varepsilon}$ does not affect cross-sectional consumption inequality. The Ramsey policymaker thus responds to such shocks using *only* monetary policy, with transfers $\boldsymbol{\tau}_x$ remaining at steady state.

5.3 General Ramsey problem

The analysis in Section 5.2 provided a sharp characterization of optimal fiscal and monetary policies in an instructive special case of our model. In future versions of this section we instead plan to present a *quantitative* evaluation, with the model parameterized to be consistent with empirical evidence on the distributional effects of monetary policy shocks (and so monetary policy in general, by the results in McKay & Wolf (2021)). In particular, we will calibrate the model to be consistent with the cross-sectional earnings and consumption response profiles documented in the prior empirical work of Guvenen et al. (2014), Coibion et al. (2017) and Amberg et al. (2021).

6 Conclusion

Should inequality affect policymakers' responses to business-cycle fluctuations? In this paper we have proposed a new method to characterize optimal policy rules in environments with rich microeconomic heterogeneity. Our method yields optimal rules in the form of forecast target criteria; appealingly, these criteria are easy to compute, robustly optimal independent of the shocks hitting the economy, and can be directly tied to empirical evidence on policy shock propagation.

This version of the draft has presented analytical results for various instructive special cases. In future versions, we plan to complement these analytical results with a quantitative model-based evaluation, with the model parameterized to be consistent with empirical evidence on the distributional effects of stabilization policy.

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.

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

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.

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

Bernanke, B. (2015). Experts Discuss Congress and Oversight of the Federal Reserve. `http://aiweb.techfak.uni-bielefeld.de/content/bworld-robot-control-software/`. [Online; accessed 21-December-2021].

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.

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.

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

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

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.

Floden, M. & Lindé, J. (2001). Idiosyncratic risk in the united states and sweden: Is there a role for government insurance? *Review of Economic Dynamics*, *4*(2), 406–437.

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

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

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.

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

Slacalek, J., Tristani, O., & Violante, G. L. (2020). Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area. *Journal of Economic Dynamics and Control*, *115*, 103879.

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.

Woodford, M. (2011). *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 presents the calibration of our illustrative Werning-type model in Section 5.2.

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

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

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

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

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

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

Integrating both sides across k yields the aggregate production (12), where $d_t \equiv \int \left(\frac{w_{kt}}{w_t}\right)^{-\varepsilon_t} dk$, 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 NKPC. We assume that union wage payments to households are subsidized at gross rate $\frac{\bar{\varepsilon}\Xi}{(\bar{\varepsilon}-1)(1-\tau_y)}$, where $\bar{\varepsilon}$ 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 C.1; 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 n_{kt} to maximize

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

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

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

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

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

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

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

Next, we from the definition of the price index have

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

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

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

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

A.2 Example model parameterization for Section 5.2

To construct Figure 1 we use the following parameterization. For household utility we set $\gamma = 1$ and $\phi = 1$. $\beta = 0.978$ is set to match a supply of liquid assets to GDP of 0.26, exactly as in Kaplan et al. (2018). The steady state real interest rate is 1% quarterly. On the supply side, we calibrate θ so that the slope of the NKPC is 0.02, consistent with the empirical evidence surveyed by Mavroeidis et al. (2014). We set the elasticity of substitution $\bar{\epsilon}$ to 6.

The idiosyncratic income process is based on Floden & Lindé (2001) with the addition of an exogenous aggregate shifter that stretches or collapses the income distribution in a mean preserving fashion. ζ_{it} is the sum of a persistent and a transitory component, with the persistence parameter 0.978 and standard deviation 0.107 of the persistent shock and standard deviation 0.045 of the transitory shock. e_{it} is then given by

$$e_{it} = \exp\left(\zeta_{it} - \min\{\zeta_{it}, 0\} m_t + \bar{e}_t\right),$$

where $\bar{e}_t$ is a normalizing factor that ensures the average $\int_0^1 e_{it} di$ is equal to 1. A positive value of the distributional shock m_t raises incomes for low values of ζ and reduces incomes at high values of ζ (through a reduction in $\bar{e}_t$).

B 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 B.1. Our computation of the inequality term in the full Ramsey loss function is described in Appendix B.2.

B.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 (22) - (23) 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" (23) leverages the so-called "fake news" algorithm.

For computation of the alternative (but equivalent) constraint formulations (25) - (27) 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.

B.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 (28), 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

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

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

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

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

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 (B.2) and (B.3):

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

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

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

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

Computing Q. To compute $\Omega_t(\zeta^t)$ and so Q, the key challenge is to arrive at the derivatives in (B.2) - (B.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 (B.4)-(B.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 (B.4), (B.5), and (B.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 (24).

C Proofs and auxiliary lemmas

C.1 Proof of Proposition 1

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:

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

Our objective is to construct a second-order approximation of (C.1). Similar to the analysis in Woodford (2011), our strategy is to consider an efficient steady state, allowing evaluation of Equation (C.1) 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

$$\int \varphi(\zeta)^{1/\gamma} d\Gamma(\zeta) = \bar{c}\bar{u}_c^{1/\gamma}. \tag{C.2}$$

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

$$\Xi \equiv \left(\int \varphi(\zeta)^{1/\gamma} d\Gamma(\zeta)\right)^{\gamma} = \varphi(\zeta)\bar{\omega}(\zeta)^{-\gamma} \qquad \forall \zeta. \tag{C.3}$$

With these preliminary definitions out of the way, we can begin constructing the second-order approximation of (C.1). 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 (C.2) 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 (C.1):

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

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 (2011). Recall that the dispersion term is defined as

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

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

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

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

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

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

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

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

Solving backwards:

$$\widehat{d}_t \approx \theta^{t+1}\widehat{d}_{-1} + \frac{\theta\bar{\varepsilon}}{2(1-\theta)}\sum_{s=0}^{t}\theta^{t-s}\widehat{\pi}_s^2 \tag{C.4}$$

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 $\frac{\theta\bar{\varepsilon}}{2(1-\theta)}(1, \beta\theta, \cdots)$ so the discounted sum is $\frac{\theta\bar{\varepsilon}}{2(1-\theta)(1-\beta\theta)}$. Using this we have

$$\begin{aligned}\sum_{t=0}^{\infty}\beta^t U_t &\approx -\nu_\ell(\bar{\ell})\bar{\ell}\sum_{t=0}^{\infty}\beta^t\left[\frac{\theta\bar{\varepsilon}}{2(1-\theta)(1-\beta\theta)}\widehat{\pi}_t^2 + \frac{1}{2}(\gamma+\phi)\,\widehat{y}_t^2 + \frac{\gamma}{2}\int\frac{\widehat{\omega}(\zeta)^2}{\bar{\omega}(\zeta)}d\Gamma(\zeta)\right] \\ &= -\frac{\nu_\ell(\bar{\ell})\bar{\ell}\theta\bar{\varepsilon}}{2(1-\theta)(1-\beta\theta)}\sum_{t=0}^{\infty}\beta^t\left[\widehat{\pi}_t^2 + \frac{\kappa}{\bar{\varepsilon}}\widehat{y}_t^2 + \frac{\kappa\gamma}{(\gamma+\phi)\,\bar{\varepsilon}}\int\frac{\widehat{\omega}(\zeta^t)^2}{\bar{\omega}(\zeta^t)}d\Gamma(\zeta^t)\right],\end{aligned} \tag{C.5}$$

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

C.2 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:

$$\Pi_\pi \widehat{\boldsymbol{\pi}} = \Pi_y \widehat{\boldsymbol{y}} + \psi \boldsymbol{\varepsilon}, \tag{C.6}$$

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

$$\widehat{\boldsymbol{y}} = \mathcal{C}_y \widehat{\boldsymbol{y}} + \mathcal{C}_r \widehat{\boldsymbol{r}} + \mathcal{C}_x \widehat{\boldsymbol{\tau}}_x + \mathcal{C}_e \widehat{\boldsymbol{\tau}}_e + \mathcal{C}_m \widehat{\boldsymbol{m}}, \tag{C.7}$$

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

$$\widehat{r}_0 = r_0(\widehat{\pi}_0, \widehat{y}_0, \widehat{q}_0) \tag{C.8}$$

$$\widehat{\boldsymbol{r}}_{+1} = r_{+1}(\widehat{\boldsymbol{i}}, \widehat{\boldsymbol{\pi}}) \tag{C.9}$$

and

$$\widehat{\boldsymbol{q}} = q(\widehat{\boldsymbol{\pi}}_{+1}, \widehat{\boldsymbol{y}}_{+1}, \widehat{\boldsymbol{r}}_{+1}) \tag{C.10}$$

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

$$\widehat{\boldsymbol{\tau}}_e = \tau_e(\widehat{\boldsymbol{y}}, \widehat{\boldsymbol{\tau}}_x, \widehat{\boldsymbol{\pi}}, \widehat{\boldsymbol{q}}). \tag{C.11}$$

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 (C.6) - (C.11) 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 (C.7); requirement 4 is re-stated in (C.6); requirement 5 is imposed by (C.11); requirement 6 is imposed by (C.8)-(C.10); 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 (C.7). First, given $\widehat{\boldsymbol{\pi}}$ and $\widehat{\boldsymbol{i}}$, we can recover $\widehat{\boldsymbol{r}}_{+1}$ from (C.9). We can thus recover $\widehat{\boldsymbol{q}}$ from (C.10), and so finally we get $\widehat{r}_0$ from (C.8). Second, given $\widehat{\boldsymbol{y}}$, $\widehat{\boldsymbol{\tau}}_x$, $\widehat{\boldsymbol{\pi}}$ and $\widehat{\boldsymbol{q}}$, we can recover $\widehat{\boldsymbol{\tau}}_e$ from (C.11). 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

$$\widehat{\boldsymbol{y}} = \underbrace{[\mathcal{C}_y + \mathcal{C}_r \mathcal{R}_y + \mathcal{C}_e \mathcal{T}_y]}_{\tilde{\mathcal{C}}_y} \widehat{\boldsymbol{y}} + \underbrace{[\mathcal{C}_r \mathcal{R}_\pi + \mathcal{C}_e \mathcal{T}_\pi]}_{\tilde{\mathcal{C}}_\pi} \widehat{\boldsymbol{\pi}} + \underbrace{[\mathcal{C}_r \mathcal{R}_i + \mathcal{C}_e \mathcal{T}_i]}_{\tilde{\mathcal{C}}_i} \widehat{\boldsymbol{i}} + \underbrace{[\mathcal{C}_x + \mathcal{C}_e \mathcal{T}_x]}_{\tilde{\mathcal{C}}_x} \widehat{\boldsymbol{\tau}}_x + \mathcal{C}_m \boldsymbol{m} \quad \text{(C.12)}$$

where $\mathcal{T}_\bullet$ and $\mathcal{R}_\bullet$ are derivative matrices for the maps $\widehat{\boldsymbol{r}}(\bullet)$ and $\widehat{\boldsymbol{\tau}}_e(\bullet)$. (C.12) embeds (13), (14) and (C.11). 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. □

C.3 Proof of Proposition 2

In light of Lemma 1, we can re-state the optimal policy problem as minimizing (18) subject to the two constraints (C.6) and (C.12). This problem gives the following necessary and sufficient first-order conditions:

$$\lambda_\pi W \widehat{\boldsymbol{\pi}} + \Pi_\pi' W \boldsymbol{\varphi}_\pi - \tilde{\mathcal{C}}_\pi' W \boldsymbol{\varphi}_y = \mathbf{0} \quad \text{(C.13)}$$

$$\lambda_y W \widehat{\boldsymbol{y}} - \Pi_y' W \boldsymbol{\varphi}_\pi + (I - \tilde{\mathcal{C}}_y') W \boldsymbol{\varphi}_y = \mathbf{0} \quad \text{(C.14)}$$

$$-\tilde{\mathcal{C}}_i' W \boldsymbol{\varphi}_y = \mathbf{0}, \quad \text{(C.15)}$$

where $\boldsymbol{\varphi}_\pi$ and $\boldsymbol{\varphi}_y$ are sequences of Lagrange multipliers on the two constraints. The proof of Proposition 2 proceeds by guessing (and then verifying) that $\boldsymbol{\varphi}_y = \mathbf{0}$. Under this assumption, we can combine (C.13) - (C.14) to get

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

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

$$\lambda_\pi \widehat{\boldsymbol{\pi}} + \frac{\lambda_y}{\kappa} \begin{pmatrix} 1 & 0 & 0 & \dots \\ -1 & 1 & 0 & \dots \\ 0 & -1 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \widehat{\boldsymbol{y}} = \mathbf{0} \tag{C.16}$$

But this is just (30), 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 (C.6) and (C.16) given $(\boldsymbol{m}, \boldsymbol{\varepsilon})$. Plugging into (C.12) and re-arranging:

$$\underbrace{\widehat{\boldsymbol{y}}^* - \tilde{\mathcal{C}}_y \widehat{\boldsymbol{y}}^* - \tilde{\mathcal{C}}_\pi \widehat{\boldsymbol{\pi}}^* - \mathcal{C}_m \boldsymbol{m}}_{\text{demand target}} = \tilde{\mathcal{C}}_i \widehat{\boldsymbol{i}} \tag{C.17}$$

It thus remains to show that condition (29) is precisely sufficient to ensure that we can always find $\widehat{\boldsymbol{i}}$ such that (C.17) holds. To see this, note that the left-hand side of (C.17) 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

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}}\right)^t \bar{y} \widehat{y}_t \tag{C.18}$$

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

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}}\right)^t \bar{c} \widehat{c}_t = \sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}}\right)^t \left\{(1+\bar{r}) \bar{a} \widehat{r}_t + (1-\tau_y) \bar{y} \widehat{y}_t + \bar{\tau}_x \widehat{\tau}_{xt} + \bar{\tau}_e \widehat{\tau}_{et}\right\} \tag{C.19}$$

Doing the same for the government budget constraint, we get

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}}\right)^t \left\{(1+\bar{r}) \bar{a} \widehat{r}_t + \bar{\tau}_x \widehat{\tau}_{xt} + \bar{\tau}_e \widehat{\tau}_{et}\right\} = \sum_{t=0}^{\infty} \tau_y \bar{y} \widehat{y}_t \tag{C.20}$$

Combining (C.20) and (C.19), we get (C.18), as claimed. This completes the argument. □

C.4 Auxiliary lemma for Proposition 3

We here establish that, under Assumptions 1 to 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 C.1. *Suppose that Assumptions 1 to 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* $t = 0, 1, 2, \ldots,$

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

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

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

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 (C.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

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

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

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

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 (C.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 (C.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:

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

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. (C.25) holds for any paths $(\boldsymbol{r},\boldsymbol{y},\boldsymbol{m},\boldsymbol{\tau}_x)$ such that (C.21) holds and $\widehat{\boldsymbol{m}} = \boldsymbol{\tau}_x = \boldsymbol{0}$. □

C.5 Proof of Proposition 3

It remains to show that changes in nominal interest rates induce paths of $\boldsymbol{r}$ and $\boldsymbol{y}$ that satisfy (C.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 1 to 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).[20] We refer the reader to Werning (2015) for the formal proof. □

C.6 Proof of Corollary 1

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

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

It follows from (C.6) that

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

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

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

giving (34). □

C.7 Proof of Corollary 2

We have already shown that the optimal monetary rule is given as (34). Since (C.6) also implies that

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

[20] 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 1 to rewrite the first two terms in (35) as

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

Imposing (34) sets these terms to zero, so Corollary 2 follows. □