On the Optimal Use of Fiscal Stimulus Payments at the Zero Lower Bound

Alisdair McKay
Federal Reserve Bank of Minneapolis

Christian K. Wolf
Massachusetts Institute of Technology
National Bureau of Economic Research

For much of the decade following the Great Recession, central banks across the world remained constrained by a binding zero (or effective) lower bound (ZLB) on nominal interest rates. Much academic and policy interest thus centered on the question of how fiscal policy could be used to manage the economy instead.[1] A key takeaway from that literature is that fiscal instruments—either unconventional (e.g., consumption and labor subsidies) or more conventional (e.g., stimulus checks)—can in principle be used to replicate monetary policy's effects on aggregate demand, thus allowing policymakers to close aggregate output gaps and stabilize inflation even at the ZLB.

While very similar in their effects on aggregate demand and thus the economy as a whole, those instruments may however differ substantially in their distributional incidence. On the one hand, interest rate policy and consumption subsidies are likely to have broad-based effects: everyone tends to benefit, and so such policies tend to be stimulative across the distribution of households.[2] On the other hand, uniform stimulus checks—as seen frequently in the U.S.

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. We thank Jordi Galí for a helpful discussion of our work.

1. See Eggertsson (2011); Christiano and others (2011); Correia and others (2013); Wolf (2021), among many others.

2. McKay and Wolf (2023a); Bachmann and others (2021).

Heterogeneity in Macroeconomics: Implications for Monetary Policy, edited by Sofía Bauducco, Andrés Fernández, and Giovanni L. Violante, Santiago, Chile.

over the past two decades—are much more progressive in their effects on household consumption: by construction, uniform transfers lead to a larger percentage change in income for low-income households—an effect that is only reinforced further by higher marginal propensities to consume at the bottom of the income and wealth distribution. If policymakers have distributional concerns in addition to their usual aggregate objectives, then this heterogeneous distributional incidence will shape optimal policy, including in particular at the ZLB.

The core contribution of this paper is to explore the optimal use of fiscal stabilization policy at the ZLB. Methodologically, doing so requires us to generalize our approach in McKay and Wolf (2023b) to environments subject to a binding ZLB constraint. Substantively, our core takeaway will be that—for canonical ZLB-type shocks, like a tightening in borrowing constraints or a distributional shock concentrated on low-income households—transfer stimulus payments are not just a mere substitute for classical unconstrained monetary policy; rather, they strictly improve upon it.

Environment. We consider a relatively standard heterogeneous-agent (HANK) model, rich enough to be consistent with the broad empirical patterns for the distributional incidence of monetary and fiscal stabilization policies. The economy is subject to a shock that disproportionately reduces the consumption of low-income households—a reduced-form stand-in for tighter borrowing constraints[3] or greater inequality.[4] The shock reduces aggregate demand and thus requires a policy response to stabilize the macroeconomy. We assume that the shock is large enough so that—in the presence of a ZLB on nominal interest rates—monetary policy alone is insufficient to stabilize aggregate demand.

Optimal Policy. We study the optimal policy problem of a policymaker that seeks to avoid cyclical changes in (i) the output gap, (ii) inflation, and (iii) the cross-sectional distribution of consumption. Such a loss function corresponds to a second-order approximation to a social welfare function where the Pareto weights are set so that the steady-state cross-sectional distribution of consumption is optimal.[5] “We refer to a policymaker with this particular loss function as the “Ramsey planner.” We ask how such a policymaker uses three available tools—standard interest rate policy, unconventional fiscal policy à

3. For example, Eggertsson and Krugman (2012); Guerrieri and Lorenzoni (2017).
4. For example, Auclert and Rognlie (2018).
5. See McKay and Wolf (2023b).

la Correia and others (2013) (i.e., consumption and labor subsidies), and uniform stimulus checks—to stabilize the economy as well as possible.[6] It will also prove instructive to contrast those results with outcomes for an alternative policymaker that only cares about output and inflation—i.e., a conventional "dual-mandate" policymaker.

Our first results concern the use of conventional nominal interest rate policy. Without the ZLB constraint, a dual-mandate policymaker would lower nominal interest rates as far as needed to perfectly close the output gap and stabilize inflation.[7] Relative to this familiar dual-mandate benchmark, our Ramsey planner would additionally like to stabilize the cross-sectional consumption distribution. However, since interest rate cuts have broad-based stimulative effects across the consumption distribution, they do little to help the planner's distributional goals. Thus, if unconstrained, the Ramsey planner would cut interest rates in a manner similar to the usual dual-mandate outcome. With a binding ZLB, this interest rate cut is of course not feasible, and so now output and inflation gaps arise in addition to the cross-sectional consumption dispersion.

We next consider the use of unconventional fiscal policy—i.e., consumption subsidies to increase consumer demand, and labor taxes to offset the labor supply effects of the consumption subsidy. A dual-mandate policymaker could use these tools to perfectly stabilize aggregate output and inflation even with a binding ZLB, as discussed by Correia and others. We find, however, that such a policy is not particularly useful to the full Ramsey policymaker: unconventional fiscal policy again stimulates consumption across the entire cross-sectional income and wealth distribution and so—just like the infeasible interest rate cut—does little to address the inequality caused by the initial shock.

Finally we turn attention to conventional fiscal policy in the form of uniform stimulus payments. Consistent with the results in Wolf (2021), such uniform stimulus checks can also be used to perfectly stabilize aggregate output and inflation. Importantly, however, stimulus payments do so largely by boosting consumption of low-income households, directly counteracting the distributional incidence

6. Our analysis therefore takes a very particular perspective on the policy problem: the goal is to offset the effects of the business cycle without changing the long-run consumption distribution.

7. This is possible in our economy by the usual "divine coincidence" argument: our economy is subject to a demand shock, and monetary policy can in principle perfectly neutralize that shock's effects on aggregates.

of the original business-cycle shock. This delivers our headline result: for a Ramsey policymaker, at a binding ZLB caused by a distributional shock mostly hitting the poor, stimulus payments do not just substitute for conventional monetary policy—they strictly improve upon it.

Literature. A vast literature has studied macroeconomic stabilization policy at the ZLB—e.g., Krugman (1998); Eggertsson and Woodford (2003); Werning (2011). Our work in particular relates to the subset of that literature that has considered the interaction of inequality and the ZLB. As mentioned briefly above, Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2017) study how deleveraging at the bottom end of the income distribution may act as a demand-type shock that pushes the economy towards the ZLB. The interaction between one classic monetary policy remedy to the ZLB—forward guidance—and inequality is analyzed in McKay and others (2016) and Farhi and Werning (2019). Closest to our focus on stimulus checks, Mehrotra (2018) and Wolf (2021) consider fiscal stimulus payments at the ZLB.

Outline. The paper is organized as follows. Section 1 introduces the HANK model and presents the optimal policy problem. The model calibration is described in section 2, and we there also discuss the distributional effects of our three policy instruments. The headline optimal policy results are presented in section 3. Section 4 concludes.

1. Model

For our optimal policy analysis, we rely on a relatively standard sticky-wage HANK model. The only nonstandard model feature is that it includes long-term bonds in addition to the usual short-term bonds. Importantly, the presence of such long-term bonds limits the extent of redistribution that occurs through changes in short-term interest rates, allowing our model to imply a realistic distributional incidence of monetary policy.

Time is discrete and runs forever, $t = 0,1,2,\ldots$ Consistent with our linear-quadratic framework in section 1.6, we will consider linearized perfect-foresight transition sequences. The perfect-foresight approach is in keeping with existing methods for analyzing business-cycle models with occasionally binding constraints on aggregate variables.[8] Throughout this section, boldface denotes time paths

8. See, e.g., Guerrieri and Iacoviello (2015); Holden (2016).

(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 $\hat{x}$.[9]

1.1 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

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

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

Household Budget. We begin with the income side of the household budget. Households are endowed with stochastic idiosyncratic labor productivity e_{it} and have total labor earnings of $(1-\tau_{\ell,t}) w_t e_{it} \ell_{it}$, where w_t is the real wage per effective unit of labor and $\tau_{\ell,t}$ is the tax rate on labor income. We let ζ_{it} be a stochastic event that determines the labor productivity of household i at date t. The event ζ_{it} itself follows a stationary Markov process, and a canonical heterogeneous agent model would set $e_{it} = \zeta_{it}$. We will instead assume that there is a function Φ that maps ζ_{it} to e_{it},

$$e_{it} = \Phi(\zeta_{it}, d_t) \text{ with } \int e_{it} di = 1 \forall d_t.$$

This mapping depends on an exogenous distributional shock, d_t, which affects the dispersion of individual labor productivities. For the analysis in section 3, this shock d_t will be the shock that moves the economy towards the ZLB—a distributional shock that mostly affects low-income households.[10] As we will describe further below, household labor supply is determined by a labor market union, so hours worked ℓ_{it} are taken as given by an individual household. Households furthermore receive a time-varying lump-sum transfer $\tau_{x,t} + \tau_{e,t} e_{it}$. Here, the first component of the transfer, $\tau_{x,t}$, is the same for all households and will be manipulated as part of the optimal policy problem. We will thus refer to it as the "exogenous" component of fiscal

9. To be precise, we use log deviations for the variables $\{y, c, \ell, 1+r, 1+i, \varepsilon\}$ and level deviations for the variables $\{\pi, \tau_x, \tau_e, \tau_\ell, \tau_c, \mathcal{R}\}$.

10. While literally modeled here as a distributional shock, it is well-known that such a shock will have very similar effects to a tightening in borrowing constraints, as e.g., considered in Guerrieri and Lorenzoni (2017).

transfers, hence the subscript x, or “fiscal stimulus payment”. The second component, $\tau_{e,t} e_{it}$, is the “endogenous” part of taxes, adjusting slowly over time to make sure that long-run fiscal budget balance is maintained.

Households use their income to consume and save. To consider unconventional fiscal stabilization policy, as in Correia and others (2013), we allow for time-varying consumption subsidies $\tau_{c,t}$. Following Correia and others, we furthermore assume that, for any given $\tau_{c,t}$, labor taxes $\tau_{\ell,t}$ adjust to offset the labor supply impact of the consumption subsidy. We thus treat $\tau_{c,t}$ as the single “unconventional” fiscal policy instrument.

Finally, households save through financial assets with expected real return r_t between periods t and $t + 1$, subject to an exogenous no-borrowing constraint. As we discuss later, households can save in multiple assets, with their returns linked by a no-arbitrage condition. In our perfect-foresight economy, all assets will earn exactly the same common realized return r_t at each date $t = 1,2....$ At date 0, however, the realized return may deviate from the ex-ante expected return, and in particular it may depend on the household’s date-0 asset composition. For simplicity we assume that portfolios have the same composition everywhere in the cross-section of households, and we let $\mathcal{R}_t$ denote the common date-t revaluation factor of household portfolios—which, again, will only be nonzero at date 0.[11]

Putting all the pieces together, the household budget constraint is

$$\frac{1}{1+r_t} a_{it} + \left(1-\tau_{c,t}\right) c_{it} = a_{it-1}\left(1+\mathcal{R}_t\right) + \left(1-\tau_{\ell,t}\right) e_{it} w_t \ell_{it} + \tau_{x,t} + \tau_{e,t} e_{it} \tag{2}$$

where a_{it} is the expected value of assets entering period $t + 1$.

Aggregate Consumption Function. The consumption-savings problem of an individual household i is to choose consumption $\boldsymbol{c}_j$ and savings $\boldsymbol{a}_i$ to maximize (1) subject to (2).

The solution is thus a mapping from paths of real wages $\boldsymbol{w}$, hours worked $\boldsymbol{\ell}$ expected real returns $\boldsymbol{r}$, transfers τ_x and τ_e, prices $\boldsymbol{p}$, shocks $\boldsymbol{d}$, and date-0 revaluation effects $\mathcal{R}_0$ to that household’s consumption $\boldsymbol{c}_i$.

11. We allow for heterogeneous household portfolios in McKay and Wolf (2023b). The conclusions of this paper are not affected by considering such a more complicated model variant.

Aggregating consumption decisions across all households, we thus obtain an aggregate consumption function $\mathcal{C}$

$$\boldsymbol{c} = \mathcal{C}\left(\boldsymbol{w}, \ell, \boldsymbol{r}, \tau_x, \tau_e, \tau_c, \tau_\ell, \boldsymbol{d}, \mathcal{R}_0\right) \tag{3}$$

Linearizing this consumption function around the deterministic steady state yields

$$\hat{\boldsymbol{c}} = \mathcal{C}_y\left(\hat{\boldsymbol{w}} + \hat{\ell}\right) + \mathcal{C}_r\hat{\boldsymbol{r}} + \mathcal{C}_{\tau x}\hat{\tau}_x + \mathcal{C}_{\tau e}\hat{\tau}_e + \mathcal{C}_{\tau c}\hat{\tau}_c + \mathcal{C}_{\tau\ell}\hat{\tau}_\ell + \mathcal{C}_d\hat{\boldsymbol{d}} + \mathcal{C}_{\mathcal{R}}\hat{R}_0, \tag{4}$$

where the derivative matrices $\mathcal{C}$ are evaluated at steady state and we have made use of the fact that it is only the product $w_t\ell_{it}$ that is relevant to the household.

1.2 Technology, Unions, and Firms

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

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 elasticity of substitution between varieties ε. 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 implies that $e_{it}\ell_{it}\frac{W_t}{p_t} = e_{it}y_t$ for all i.

Each union sets its nominal wage in the usual 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 preferences additively separable, it follows that all households share a common marginal disutility of labor. The marginal utility of consumption, however, need not be equalized. Following McKay and Wolf (2023b)—and similar to Auclert and others (2021)—, we assume that the union evaluates the benefits of higher after-tax income using the marginal utility of average consumption $\left(c_t^{-\gamma}\right)$ rather than the average of individual household marginal utilities $\left(\int_0^1 c_{it}^{-\gamma} di\right)$. This assumption eliminates the impact of inequality on the supply side of the economy, and so we overall arrive at the following standard linearized perfect-foresight New Keynesian Phillips Curve (NKPC):

$$\hat{\pi}_t = \kappa \hat{y}_t + \beta \hat{\pi}_{t+1} \tag{5}$$

where $\kappa \equiv (\phi + \gamma)\frac{(1-\theta)(1-\beta\theta)}{\theta}$. In our derivation of (5), we allow for a (time-invariant) subsidy on union labor hiring, financed with lump-sum taxes also levied on the unions; this subsidy will yield the efficiency of the deterministic steady state needed to specify our optimal policy problem in a form consistent with a linear-quadratic analysis, as in Woodford (2003).

1.3 Asset Structure

There are two different assets in the economy: a short-term, nominal bond in zero net supply, and a long-term bond in positive net supply. By a no-arbitrage condition, both assets will provide the same expected returns along equilibrium transition paths (except possibly at $t = 0$), thus allowing us to consider a single asset in the household budget constraint (2). The realized return at date 0, however, will generally differ between the two assets. As mentioned above, the purpose of the long-term bond is to provide a more realistic description of the passthrough of monetary policy to household interest payments.

A unit of the short-term bond purchased at time t then returns $\frac{1+i_t}{1+\pi_{t+1}}$ units of the final good at time $t + 1$. For the long-term bond, at time t, households can purchase a unit of the bond for a real price of q_t (i.e., denominated in goods); at time $t + 1$, the household then receives a real "coupon" of $(\bar{r}+\delta)(1+\pi_{t+1})^{-1}$ and furthermore retains a fraction $(1-\delta)(1+\pi_{t+1})^{-1}$ of the initial asset position, now valued at $(1-\delta)(1+\pi_{t+1})^{-1} q_{t+1}$ in units of goods. Note that the parameter δ controls the duration of the bond, with lower values of δ corresponding to higher duration. The coupon scaling factor $(\bar{r} + \delta)$ is chosen to normalize the steady-state price of the bond to one. Finally, the presence of the inflation terms reflects the fact that the bond is nominal, so inflation reduces the real value of the current and future coupons, and so reduces the real value of the bond position.

Overall, it follows that the price of the long-term bond satisfies

$$q_t = \frac{(\bar{r}+\delta)(1+\pi_{t+1})^{-1} + (1-\delta)(1+\pi_{t+1})^{-1} q_{t+1}}{1+r_t}, \tag{6}$$

where r_t is the real interest rate between t and $t+1$. Real returns are furthermore linked to returns on the short-term rate via the standard Fisher relation

$$1+r_t = \frac{1+i_t}{1+\pi_{t+1}}. \tag{7}$$

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 bond, so time-0 realized returns are simply those on the long-term bond, which implies that

$$1+\mathcal{R}_0 = \frac{(\bar{r}+\delta)(1+\pi_0)^{-1} + (1-\delta)(1+\pi_0)^{-1} q_0}{1+\bar{r}}. \tag{8}$$

Note that this relation expresses the scaling factor $1+\mathcal{R}_0$ as the ratio of the actual realized return on the long-term bond (i.e., the numerator) to the expected return (i.e., the steady-state real rate in the denominator).

1.4 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-term bond. Letting $a_t^g(1+\mathcal{R}_0)$ denote the value of claims on the government entering period t (inclusive of returns), the government budget constraint becomes

$$\frac{a_{t+1}^g}{1+r_t} = a_t^g(1+\mathcal{R}_t) + \tau_{x,t} + \tau_{e,t} + \tau_{c,t}c_t - \tau_{\ell,t}y_t. \tag{9}$$

Note that, when news arrives, the claims on the government are revalued in the exactly same way as previously discussed for the household sector's assets.

We consider the nominal rate of interest, i_t, the exogenous component of transfers, $\tau_{x,t}$, and the consumption subsidy, $\tau_{c,t}$, as the independent policy instruments of the government, used for business-cycle stabilization policy. The time-varying labor tax furthermore adjusts automatically so that the net consumption benefit of an hour of work is unaffected by the consumption subsidy, requiring that $1 - \tau_{\ell,t}$ be proportional to $1 - \tau_{c,t}$ at all times. Since all three policy instruments will generally have budgetary implications, it remains to specify how long-term budget balance is ensured. We will assume that the endogenous component of transfers $\tau_{e,t}$ adjusts gradually according to the rule

$$\tau_{e,t} = \left(\bar{r} + \delta\right)\left(a_t^g - \bar{a}^g\right) \tag{10}$$

where $\bar{a}^g$ is the real, steady-state value of government debt.

1.5 Equilibrium

Given paths of exogenous shocks $\{d_t\}_{t=0}^{\infty}$ and policy instruments $\{i_t, \tau_{x,t}, \tau_{c,t}\}_{t=0}^{\infty}$, a perfect-foresight equilibrium of our linearized economy is a set of sequences of endogenous aggregate variables $\{a_t^g, c_t, y_t, q_t, a_t, \pi_t, r_t, \ell_t, \tau_{e,t}, \tau_{\ell,t}\}_{t=0}^{\infty}$ and $\mathcal{R}_0$ that satisfy the following conditions:

1. The path of aggregate consumption $\{c_t\}_{t=0}^{\infty}$ is consistent with the linearized aggregate consumption function (4), and the path of household asset holdings $\{a_t\}_{t=0}^{\infty}$ is consistent with the budget constraint (2), aggregated across households.

2. The paths of $\{\ell_t, y_t\}_{t=0}^{\infty}$ satisfy the linearized aggregate production function $y_t = \ell_t$.

3. The paths $\{\pi_t, y_t\}_{t=0}^{\infty}$ are consistent with the Phillips curve (5).

4. The evolution of government debt a_t^g, the endogenous component of transfers $\tau_{e,t}$, and the labor income tax $\tau_{\ell,t}$ are consistent with the budget constraint (9), the law of motion (10), and the requirement that $\left(1 - \tau_{\ell,t}\right) / \left(1 - \tau_{c,t}\right) = \left(1 - \bar{\tau}_\ell\right) / \left(1 - \bar{\tau}_c\right)$.

5. The asset prices $\{q_t, r_t\}_{t=0}^{\infty}$ satisfy (6) and (7), and the revaluation effect $\mathcal{R}_0$ satisfies (8).

6. The output and asset markets clear, so $y_t = c_t$ and $a_t = a_t^g$.

Note that this definition of equilibrium takes the policy instruments $\{i_t, \tau_{x,t}, \tau_{c,t}\}_{t=0}^{\infty}$ as given. The paths for these will be determined by solving the optimal policy problem.

1.6 The Policy Problem

We consider a policymaker who wishes to both stabilize the aggregate economy and offset cyclical changes in consumption inequality.

Objective Function. To understand our formulation of the policymaker's objective, we begin by noting that households in our model are ex ante identical and only differ ex post due to different realizations of their idiosyncratic shocks. Households can therefore be indexed by the history of idiosyncratic shocks they have experienced, denoted $\zeta_i^t \equiv (\zeta_{it}, \zeta_{it-1}, \ldots)$. As the shocks ζ_{it} are drawn from a stationary process, the distribution of such histories is itself stationary. With this notation established, we write the policymaker objective as

$$\mathcal{L} \equiv \sum_{t=0}^{\infty} \beta^t \left[\hat{\pi}_t^2 + \frac{\kappa}{\varepsilon} \hat{y}_t^2 + \frac{\kappa}{\varepsilon} \frac{\gamma}{\gamma + \phi} \int \frac{\hat{\omega}_t(\zeta)^2}{\overline{\omega}(\zeta)} d\Gamma(\zeta) \right] \tag{11}$$

where ζ is an infinite history of idiosyncratic shocks, $\omega_t\left(\zeta_i^t\right) \equiv c_{it} / c_t$ is the consumption share of an individual with that history at date, t and Γ is the stationary distribution of household idiosyncratic histories. In McKay and Wolf (2023b), working with a very similar model, we derive the loss function (11) as a second-order approximation to a particular social welfare function—one that attaches Pareto weights to the welfare of individual households in exactly the right way to ensure that the policymaker does not wish to deviate from the steady-state distribution of household consumption.

Next, since household consumption and thus in particular the consumption shares $\omega_t\left(\zeta_i^t\right)$ are a function solely of the aggregate variables that influence the household's consumption-savings problem, it is straightforward to see that we can re-write (11) as[12]

12. A detailed argument—including details on how to compute Q—are provided in McKay and Wolf (2023b).

$\mathcal{L} = \hat{x}' Q \hat{x},$

where Q is a symmetric matrix and $\boldsymbol{x}$ stacks the paths of the various endogenous and exogenous variables entering the consumer problem,

$x = \left(\boldsymbol{y}' \ \boldsymbol{\pi}' \ \boldsymbol{r}' \ \tau_e' \ \mathcal{R}' \ \boldsymbol{i}' \ \tau_x' \ \tau_c' \ \boldsymbol{d}'\right)'.$

Constraints. We now turn to the constraints on the policy problem. Using sequence-space methods, we can compactly express the equilibrium of this economy as

$$\hat{\boldsymbol{x}} = \underbrace{\Theta_d \hat{\boldsymbol{d}}}_{\equiv \bar{\boldsymbol{x}}} + \Theta_i \hat{\boldsymbol{i}} + \Theta_x \hat{\tau}_x + \Theta_c \hat{\tau}_c \tag{12}$$

where the Θ's are general equilibrium impulse response matrices to the shock $\hat{d}$ and the policy instruments $\{\hat{\boldsymbol{i}}_b, \hat{\tau}_x, \hat{\tau}_c\}$; and $\bar{\boldsymbol{x}}$ denotes outcomes if the policy instruments were not adjusted in response to the shock $\hat{\boldsymbol{d}}$.[13]

In addition, policy is constrained by a lower bound on the nominal interest rate. As we work with the model in deviations from a zero-inflation steady state, we express the ZLB constraint as $\hat{i} \geq \underline{i} = -\bar{r}$. We impose no constraints on the other two policy instruments.

Policy Problem. We can express the policy problem compactly by defining $\boldsymbol{p} \equiv \left(\hat{\boldsymbol{\imath}}, \hat{\tau}_x', \hat{\tau}_c'\right)'$ as the vector of policy instruments, letting $\underline{\boldsymbol{p}}$ denote the lower bounds on the instruments (which are $-\infty$ for the two fiscal instruments), and finally defining $\Theta_p \equiv \left(\Theta_i, \Theta_x, \Theta_c\right)$. We then solve the problem

$$\min_{p,\hat{x}} \hat{x}' Q \hat{x} \tag{13}$$

subject to

$$\hat{\boldsymbol{x}} = \bar{\boldsymbol{x}} + \Theta_p \boldsymbol{p} \tag{14}$$

$$\boldsymbol{p} \geq \underline{\boldsymbol{p}}. \tag{15}$$

13. In practice, to compute the Θ's, we truncate the transition paths at some large (but finite) horizon T and assume the economy has returned to steady state by this time. As there are nine variables in x, each $\Theta_\bullet$ is a $9T \times T$ matrix. See McKay and Wolf (2023b) for a discussion of how the Θ.'s are defined uniquely through policy shocks to a given baseline, determinacy-inducing monetary policy rule.

The policy problem therefore fits into a standard quadratic programming form.[14]

Finally, for reference, we will also find it useful to solve a simplified version of this problem for a dual-mandate policymaker—i.e., a policymaker with preferences as in (11), but ignoring the inequality-related term. This problem fits into (13) for a different (simpler) Q.

2. Model Parameterization

This section presents the model parameterization used for our analysis in section 3. We first discuss the calibration strategy in section 2.1 and then in section 2.2 focus on the model feature that matters most for our later results—the distributional incidence of policy.

2.1. Calibration Strategy

We provide a relatively brief sketch of our calibration strategy. A summary of the calibration is provided in table 1.

Table 1. Model Calibration

Parameter	*Description*	*Value*	*Calibration target*
γ	Relative risk aversion	1.2	Monetary shock effects
ϕ	Frisch elasticity	1	Standard
β	Discount factor	0.987	Asset market clearing
κ	Phillips curve slope	0.022	Monetary shock effects
ε	Labor Substitutability	6	Basu & Fernald (1997)
δ	Long-term bond duration	0.025	10-year maturity

Source: Authors' calculations.

14. For our numerical applications, we have found that guessing and verifying a horizon n over which the ZLB is binding to be a reliable computational strategy. In particular, given a candidate value of n, we first solve the simpler sub-problem in which the constraint binds as an equality constraint for n periods. We then verify the guess ex post. Appendix B provides details.

Households. We begin with preferences. We set the coefficient of relative risk aversion to 1.2, allowing us to match the empirically measured sensitivity of aggregate consumption to monetary policy shocks. The elasticity of labor supply is set to a standard value of 1. Next, the discount factor β is calibrated to match the steady-state level of aggregate assets in the economy. We set this asset supply to 1.4 times GDP, as in McKay and others (2016), with the implicit interpretation that assets in our model correspond to liquid assets. Turning to the idiosyncratic income process, we associate ζ_{it} with the persistent AR(1) process in the estimates of Floden and Lindé (2001) adapted to a quarterly frequency, which results in a persistence of 0.978 and an innovation variance of 0.0114. The function Φ is then given by

$$\log e_{it} = \log \zeta_{it} (1 + d_t) - \bar{e}_t,$$

where d_t is the exogenous distributional shock with $\bar{d} = 0$, and $\bar{e}_t$ is a normalization constant so that the cross-sectional average of e_{it} is always 1. Notice that an increase in d_t amplifies the dispersion in labor productivity by amplifying the differences in ζ_{it}—that is, it is an inequality shock that redistributes from the poor to the rich.

Assets and Government. We assume that households save in long-term bonds with a maturity of ten years, which corresponds to $\delta = 0.025$. The steady-state real interest rate is set to 2.4 percent per annum. Steady-state consumption subsidies are zero, and the steady-state tax rate on labor income $\bar{\tau}_\ell$ is then determined endogenously to satisfy the government budget constraint.

Supply Block. We calibrate the slope of the Phillips curve to 0.022 in order to match the magnitude of the response of inflation to a monetary policy shock.[15] Finally, the elasticity of substitution between varieties of intermediate goods is set to 6, following Basu and Fernald (1997).

2.2 The Distributional Implications of Policy

As established in prior work,[16] the three policy instruments available to our policymaker are equivalent in their effects on macroeconomic aggregates—they all equally flexibly perturb aggregate net excess demand. For optimal Ramsey policy, however, their distributional effects

15. See McKay and Wolf (2023b).
16. Correia and others (2013); Wolf (2021).

also matter. We here show that stimulus payments have very different distributional incidence from standard monetary and unconventional fiscal policy. The results are displayed in figure 1.

Monetary Policy. The top-left panel of figure 1 reveals that monetary policy in our model is broadly distributionally neutral: an interest rate cut stimulates consumption across the entire wealth distribution. This feature of our model is broadly consistent with prior empirical evidence on the heterogeneous effects of monetary policy, as for example reviewed in McKay and Wolf (2023a). We conclude that monetary policy is unlikely to have a material impact on the inequality term in the policymaker loss function (11).[17]

Figure 1. Effects of Policy Instruments on Consumption Inequality

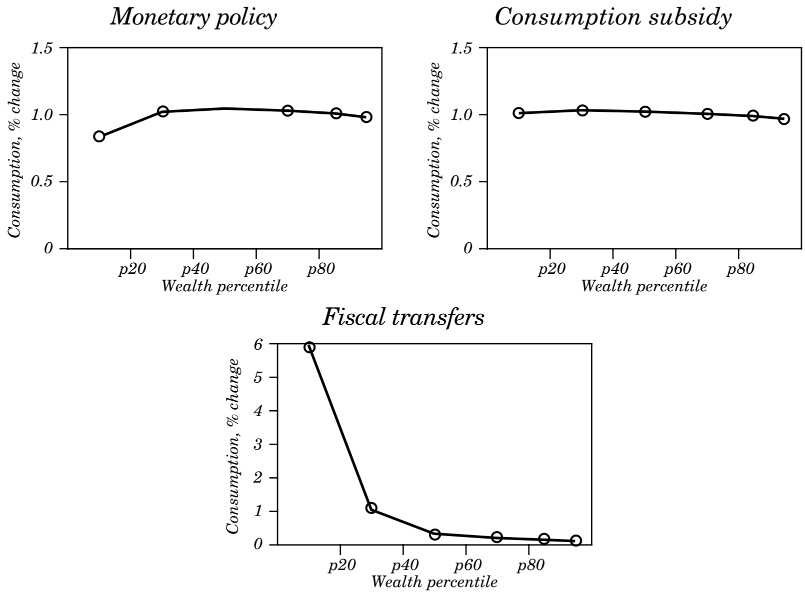

Source: Authors' calculations.

Note: The figures show the initial ($t = 0$) change in consumption following a policy stimulus. In the top-left panel, we consider an expansionary monetary policy shock that induces a one-percent increase in aggregate output on impact. Thereafter, aggregate output decays with a persistence of 0.7. In the top-right panel, we consider an unconventional fiscal policy and, in the bottom panel, we study transfer payments, with both chosen to induce a response of output of the same size and persistence as the monetary policy shock.

17. More precisely, if monetary policy were to move all households exactly up and down in tandem (and at all horizons), then monetary policy would not affect consumption shares at all, and so Ramsey policy would be identical to dual-mandate policy.

Unconventional Fiscal Policy. The top-right panel of figure 1 shows the response of consumption to a consumption subsidy, where that subsidy is chosen to replicate the aggregate output effects of the conventional monetary policy studied previously. We see that the effects on inequality are very similar to the nominal interest rate cut: households across the entire net worth distribution increase their consumption in response to the consumption subsidy. Empirically, this feature of our model is consistent with prior evidence.[18] Theoretically, the close agreement between the top-left and top-right panels follows the arguments in Seidl and Seyrich (2023).

Stimulus Checks. Finally, the bottom panel of figure 1 shows how the cross-section of consumption responds to a stimulus payment policy. The stimulative effects on consumption are now not broad-based: the consumption of the poor is disproportionately stimulated, mainly reflecting (i) the fact that the initial transfer is a larger share of their steady-state income, and (ii) their higher marginal propensities to consume. At the top end of the distribution, consumption rises mainly because, in general equilibrium, higher inflation leads to a decline in real rates, thus inducing intertemporal substitution. The differences in distributional incidence across the three policy instruments documented in figure 1 will be key to understand our optimal policy results in the next section.

3. Optimal Policy Results

This section presents our headline results on optimal stabilization policy at the ZLB. We proceed in three steps, with one subsection for each of the three policy instruments: monetary policy in section 3.1, unconventional fiscal policy in section 3.2, and fiscal stimulus payments in section 3.3. Throughout we consider an economy subject to a contractionary distributional demand shock d_t, where that shock is large enough so that the ZLB constraint becomes binding for conventional monetary policy.

3.1 Monetary Policy

We begin with conventional nominal interest rate policy. Results for optimal Ramsey monetary policy subject to the ZLB are displayed

18. Bachmann and others (2021).

as the solid lines in figure 2. By construction, the inequality shock is sufficiently large so that, given the ZLB constraint, monetary policy is unable to stabilize aggregate output. We see that nominal interest rates are cut as much as possible and remain at zero for a couple of quarters (top right), leading to an initial decline and then an overshoot in output (bottom left). The overshooting of output reflects the usual "low-for-longer" logic of optimal monetary policy at the ZLB.[19] Perhaps most importantly, monetary policy fails to counteract the distributional implications of the shock—consumption drops sharply for low-income households while remaining relatively stable for the rich (bottom right).

Figure 2. Optimal Monetary and Unconventional Fiscal Policies

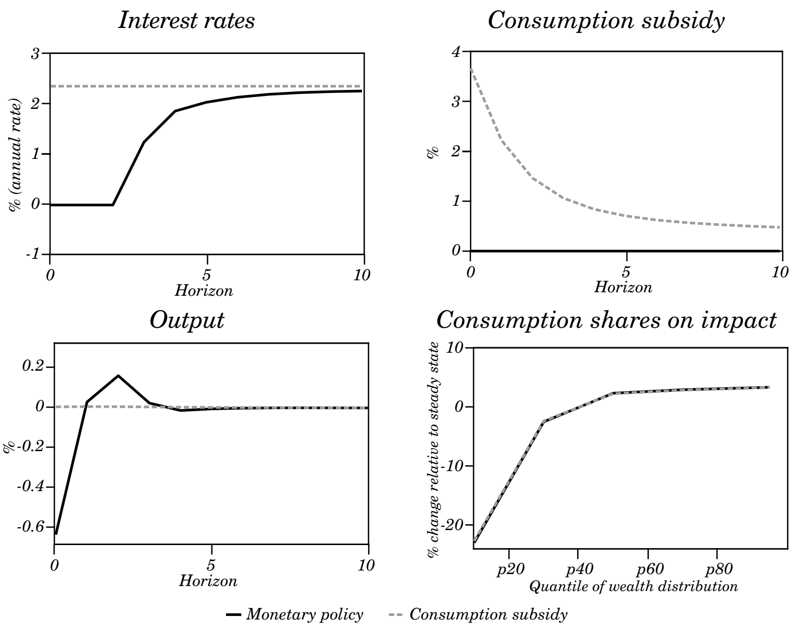

Source: Authors' calculations.

Note: Impulse responses of nominal interest rates, the consumption subsidy, output, and consumption shares (at $t = 0$) to the inequality shock under optimal (Ramsey) monetary policy subject to the ZLB (solid line) and unconventional fiscal policy (dashed).

19. See, e.g., Eggertsson and Woodford (2003).

In Appendix C we show what happens, first, in the counterfactual absence of a ZLB constraint, and second, under optimal dual-mandate monetary policy. Naturally, optimal Ramsey policy without the ZLB constraint would lower interest rates more aggressively, thus stabilizing output. Importantly, however, this additional interest rate cut does little to counteract the distributional implications of the initial shock, so consumption shares still decline significantly at the bottom end of the income and wealth distribution. This reflects the same logic as figure 1—monetary policy has small effects on the shape of the consumption distribution. In light of this, it is furthermore also not surprising that the optimal Ramsey policy looks rather similar to optimal dual-mandate policy. As monetary policy has relatively little power to moderate the effects of the initial demand shock on inequality, even the Ramsey policy is essentially only concerned with aggregate stabilization.

3.2 Unconventional Fiscal Policy

We next consider unconventional fiscal policy, as analyzed in Correia and others (2013). Results for the optimal Ramsey unconventional fiscal policy are displayed as the dashed lines in figure 2. The top-right panel shows that, as expected, the policymaker finds it optimal to subsidize consumption, thus spurring aggregate demand and almost perfectly stabilizing the macro-economy as a whole. However, as we see in the bottom-right panel, this policy does relatively little to offset the distributional tilt of the original inequality shock, with consumption shares of low-income households still dropping substantially. This is again exactly what was expected in light of figure 1: unconventional fiscal policy has broad-based stimulative effects, and so—just like conventional monetary policy—it is relatively ill-suited to deal with explicitly distributional shocks.

We note that our numerical findings in figure 2 are consistent with analytical results in Seidl and Seyrich (2023). These authors show that, for a particular mix of unconventional fiscal policy and government debt issuance, household-by-household outcomes are exactly identical to monetary stimulus. In our case the equivalence is not exact (as we consider a somewhat different debt issuance policy), but the broad intuition remains: both interest rate and consumption subsidy policy affect households in similar ways, and in particular—at least in our model calibration—those effects are rather uniform cross-sectionally.

3.3 Stimulus Payments

Finally we consider our main policy alternative: uniform stimulus payments. Results for optimal Ramsey stimulus transfer policy are displayed as solid lines in figure 3. We see that the policymaker decides to optimally pay out a relatively large positive transfer (top right), thus almost perfectly stabilizing—in fact even slightly overshooting—aggregate output (bottom left). Crucially, however, and differently from the results for monetary and unconventional fiscal policy in figure 2, this stimulus to demand at the same time also stabilizes the cross-sectional consumption distribution (bottom right). Intuitively, stimulus checks increase aggregate demand precisely by boosting the spending of those households that were hardest hit by the initial contractionary demand shock. From the point of view of our Ramsey planner with loss function (11), such transfer payments are thus the ideal tool: with one instrument, they can stabilize all three terms of their loss function. Indeed, as we show in Appendix C, the Ramsey loss for a policymaker that only uses stimulus checks is an order of magnitude lower than the loss for a counterfactual monetary policy Ramsey planner, even without the ZLB constraint. Transfers thus do not merely substitute for—but in fact very much improve upon—conventional monetary policy, at least in response to the kind of distributional demand shock that we consider here.[20]

For comparison, the dashed lines in figure 3 display optimal joint monetary-fiscal policy, which sets both stimulus payments τ_x as well as interest rates $\boldsymbol{i}$ optimally. To stabilize cross-sectional consumption shares, the stimulus payment-only policy induced a slightly excessive aggregate demand boom, overheating the economy. The optimal joint policy thus features a very mild increase in interest rates, thus closing the output gap while also keeping consumption shares stable. Overall, however, the difference in loss between transfer-only and joint optimal policy is minimal, in particular compared to the large losses that occurred under monetary-only or unconventional fiscal policy.

20. Naturally, for more broad-based initial demand shocks, conventional monetary policy or unconventional fiscal policy would again emerge as more suitable tools. However, we view explicitly distributional shocks as relevant empirically: tightening borrowing constraints were particularly important in the Great Recession, and the COVID-19 recession mostly impacted the bottom of the income distribution.

Figure 3. Optimal Stimulus Payments

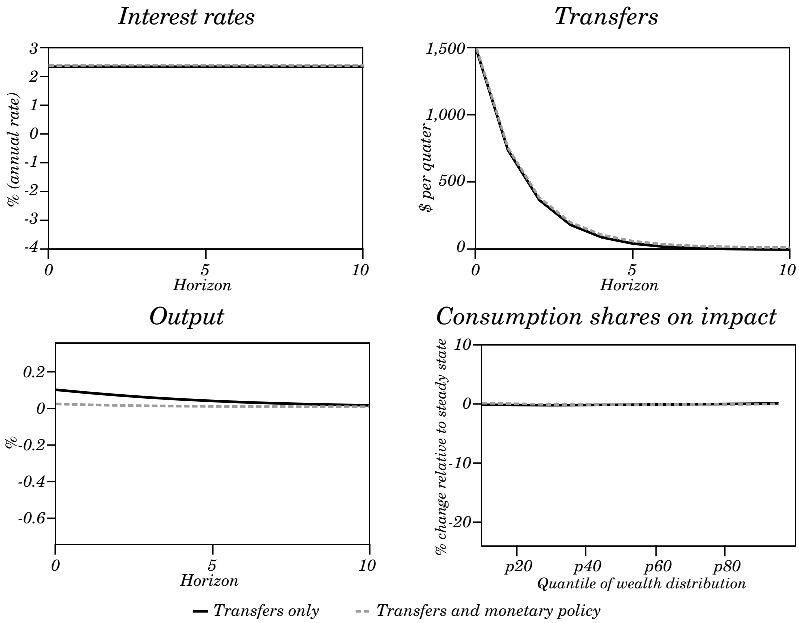

Source: Authors' calculations.

Note: Impulse responses of nominal interest, uniform stimulus payments, output, and consumption shares (at $t = 0$) to the inequality shock under optimal (Ramsey) stimulus payment policy (solid) and optimal joint monetary-transfer policy (dashed).

4. Conclusion

How should policymakers stabilize the macro-economy when conventional monetary policy is constrained by a zero lower bound on nominal rates? In particular, what policy options are most attractive if—as seems empirically plausible—the economy was subject to a negative shock that mostly impacted low-income households? Building on our prior work in McKay and Wolf (2023b), we here tried to answer those questions through the lens of a textbook heterogeneous-household model. Our headline result was that stimulus checks are more than a substitute for conventional monetary policy; in fact, since they are much more well-adapted to the distributional incidence of the shock, they are strictly preferable as a tool for cyclical stabilization, and so the ZLB poses no meaningful constraint on the policymaker.

We emphasize two important qualifiers of our results. First, our findings are necessarily sensitive to a key feature of our model—the distributional neutrality of monetary policy interventions. While this model feature is consistent with prior work,[21] further empirical investigation would be welcome. Second, our conclusions apply to particular, empirically relevant kinds of demand shocks—those mostly affecting low-income households. Conclusions may be different for other types of demand disturbances, e.g., those directly affecting firms rather than households.

21. See, e.g., McKay and Wolf (2023a) and the references therein.

References

Auclert, A. and M. Rognlie. 2018. "Inequality and Aggregate Demand." Technical Report, National Bureau of Economic Research.

Auclert, A., M. Rognlie, M. Souchier, and L. Straub. 2021. "Exchange Rates and Monetary Policy with Heterogeneous Agents: Sizing up the Real Income Channel." Technical Report, National Bureau of Economic Research.

Bachmann, R., B. Born, O. Goldfayn-Frank, G. Kocharkov, R. Luetticke, and M. Weber. 2021. "A Temporary VAT Cut as Unconventional Fiscal Policy." Technical Report, National Bureau of Economic Research.

Basu, S. and J.G. Fernald. 1997. "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy* 105(2): 249–83.

Christiano, L., M. Eichenbaum, and S. Rebelo. 2011. "When is the Government Spending Multiplier Large?" *Journal of Political Economy* 119(1): 78–121.

Correia, I., E. Farhi, J.P. Nicolini, and P. Teles, P. 2013. "Unconventional Fiscal Policy at the Zero Bound." *American Economic Review* 103(4): 1172–211.

Eggertsson, G.B. 2011. "What Fiscal Policy is Effective at Zero Interest Rates?" *NBER Macroeconomics Annual* 25: 59–112.

Eggertsson, G.B. and P. Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics* 127(3): 1469–513.

Eggertsson, G.B. and M. Woodford. 2003. "Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* (1): 139–233.

Farhi, E. and I. Werning. 2019. "Monetary Policy, Bounded Rationality, and Incomplete Markets." *American Economic Review* 109(11): 3887–928.

Floden, M. and J. Lindé. 2001. "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" *Review of Economic Dynamics* 4(2): 406–37.

Guerrieri, L. and M. Iacoviello. 2015. "Occbin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily." *Journal of Monetary Economics* 70: 22–38.

Guerrieri, V. and G. Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *Quarterly Journal of Economics* 132(3): 1427–67.

Holden, T.D. 2016. "Computation of Solutions to Dynamic Models with Occasionally Binding Constraints." EconStor Preprints No. 144569.

Krugman, P.R. 1998. "It's Baaack!: Japan's Slump and the Return of the Liquidity Trap." *Brookings Papers on Economic Activity* 29(2): 137–206.

McKay, A., E. Nakamura, and J. Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review* 106(10): 3133–58.

McKay, A. and C.K. Wolf. 2023a. "Monetary Policy and Inequality." *Journal of Economic Perspectives* 37(1): 121–44.

McKay, A. and C.K. Wolf. 2023b. "Optimal Policy Rules in HANK." Working Paper, FRB Minneapolis.

Mehrotra, N. R. 2018. "Fiscal Policy Stabilization: Purchases or Transfers?" *International Journal of Central Banking* 14(2): 1–50.

Seidl, H. and F. Seyrich. 2023. "Unconventional Fiscal Policy in HANK." DIW Berlin Discussion Paper.

Werning, I. 2011. "Managing a Liquidity Trap: Monetary and Fiscal Policy." Technical Report, National Bureau of Economic Research.

Wolf, C.K. 2021. "Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result." Working Paper No. 29193, National Bureau of Economic Research.

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

APPENDICES

Appendix A. Supply Side and Phillips Curve Derivation

We here provide further details for the production side of our economy, as sketched in section 1.2. We begin by specifying the details of the economy's production technology and then derive our Phillips curve (5).

Technology. A unit continuum of unions, indexed by $k \in [0,1]$, differentiates 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

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

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

$$W_t = \left(\int W_{kt}^{1-\varepsilon} dk \right)^{1/(1-\varepsilon)},$$

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

$$\ell_{kt} = \left(\frac{W_{kt}}{W_t} \right)^{-\varepsilon} y_t. \tag{A.1}$$

Integrating both sides across k yields the aggregate production

$$y_t \int \left(\frac{W_{kt}}{W_t} \right)^{-\varepsilon} dk = \ell_t,$$

with ℓ_t denoting total effective hours supplied by households and the integral term capturing the efficiency losses due to price dispersion. The dispersion term disappears in a first-order approximation to the dynamics of the model.

From Union Problem to NKPC. We assume that union wage payments to households are subsidized at gross rate $\frac{\varepsilon}{(\varepsilon-1)(1-\bar{\tau}_\ell)}$, where $\bar{\tau}_\ell$ is the steady-state labor income tax. The union's problem is to choose the reset wage W^* and ℓ_{kt} to maximize

$$\sum_{s\geq 0}\beta^s\theta^s\left[u_c\left(c_{t+s}\right)\frac{1-\tau_{\ell,t}}{1-\tau_{c,t}}\frac{\varepsilon}{\left(\varepsilon-1\right)\left(1-\bar{\tau}_\ell\right)}\frac{W^*}{p_{t+s}}\ell_{kt}-\nu_\ell\left(\ell_{t+s}\right)\ell_{kt}\right]$$

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

$$\sum_{s\geq 0}\beta^s\theta^s\nu_\ell\left(\ell_{t+s}\right)y_{t+s}\epsilon_{t+s}\left(\frac{p_{t+s}}{p_t}\right)^{\varepsilon}=\frac{\varepsilon}{\varepsilon-1}\sum_{s\geq 0}\beta^s\theta^s u_c\left(c_{t+s}\right)\left(\varepsilon-1\right)\frac{W_t^*}{p_t}\left(\frac{p_{t+s}}{p_t}\right)^{\varepsilon-1}y_{t+s}, \quad \text{(A.2)}$$

where w_t^* is the optimal reset wage chosen at date t, and we have used the fact that $(1-\tau_{\ell,t})\,/\,(1-\tau_{c,t})$ is constant and equal to $1-\bar{\tau}_\ell$. Log-linearizing the first-order condition around a zero-inflation steady state:

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

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

$$\hat{W}_t^*-\hat{p}_t=\left(1-\beta\theta\right)\sum_{s\geq 0}\beta^s\theta^s\left(\left(\phi+\gamma\right)\hat{y}_{t+s}+\hat{p}_{t+s}-\hat{p}_t\right).$$

Next, from the definition of the price index, we have

$$1+\pi_t\equiv\frac{p_t}{p_{t-1}}=\left(\theta^{-1}-\frac{1-\theta}{\theta}\left(\frac{W_t^*}{p_t}\right)^{1-\varepsilon}\right)^{\frac{1}{\varepsilon-1}}. \quad \text{(A.3)}$$

Log-linearizing around a zero inflation steady state, this gives

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

Eliminating $\hat{W}_t^*-\hat{p}_t$ and simplifying, we get

$$\hat{\pi}_t=\kappa\hat{y}_t+\beta\hat{\pi}_{t+1},$$

where $\kappa=\frac{\left(1-\theta\right)\left(1-\beta\theta\right)\left(\phi+\gamma\right)}{\theta}$.

Appendix B. Computation of Optimal Policy

Here we will describe how to solve the problem for an application for which the ZLB binds for the first n periods after the shock occurs. We partition $\boldsymbol{p} = (\boldsymbol{p}_1' \, \boldsymbol{p}_2')'$, where the lower bound binds on $\boldsymbol{p}_1$ and not on $\boldsymbol{p}_2$. We can then rewrite the policy problem as

$$\min_{p_2, \hat{x}} \frac{1}{2} \hat{\boldsymbol{x}}' \boldsymbol{Q} \hat{\boldsymbol{x}}$$
$$\text{s.t.} \hat{\boldsymbol{x}} = \bar{\boldsymbol{x}} + \Theta_{p,1} \underline{\boldsymbol{p}}_1 + \Theta_{p,2} \boldsymbol{p}_2,$$

where we have partitioned Θp to correspond to the partition of $\boldsymbol{p}$. The first-order conditions of this problem yield

$$\Theta_{p,2}' \boldsymbol{Q} \left(\bar{\boldsymbol{x}} + \Theta_{p,1} \underline{\boldsymbol{p}}_1 + \Theta_{p,2} \boldsymbol{p}_2 \right) = 0,$$

which we can easily solve for $\boldsymbol{p}_2$.

To solve the full problem, we perform the above calculation for all possible values of n (within reason). For each candidate n, we solve for $\boldsymbol{p}_2$ as above and then check if it violates the constraint $\boldsymbol{p}$. If so, we discard this candidate. If not, we compute and store the objective value $\hat{\boldsymbol{x}}' \boldsymbol{Q} \hat{\boldsymbol{x}}$. After evaluating all the candidate values of n, we select the one that yields the lowest objective value.

This procedure is a simple and robust method for typical macroeconomic shocks that mean revert, resulting in a binding ZLB only for the first n periods. For more complicated ZLB episodes, one could use more sophisticated quadratic programming methods.

Appendix C. Further Optimal Policy Results

This appendix presents two additional sets of optimal policy results. First, figure C.1 shows optimal monetary policy for the dual-mandate policymaker and in the absence of the ZLB constraint. Second, table C.1 shows the loss function values achieved by policymakers using different policy tools.

Table C.1 Ramsey Loss Achieved Relative to Monetary Policy

Policy instrument	*Relative loss*
Monetary policy (ZLB)	1.00
Monetary policy (unconstrained)	0.99
Unconventional fiscal policy	0.95
Fiscal stimulus payments	0.05
Joint monetary-transfer policy	0.04

Source: Authors' calculations.

Note: The table shows the policymaker loss under optimal policy for different policy tools. All values are reported relative to the loss achieved by ZLB-constrained Ramsey monetary policy.

Figure C.1 Optimal Ramsey and Dual-Mandate Monetary Policies

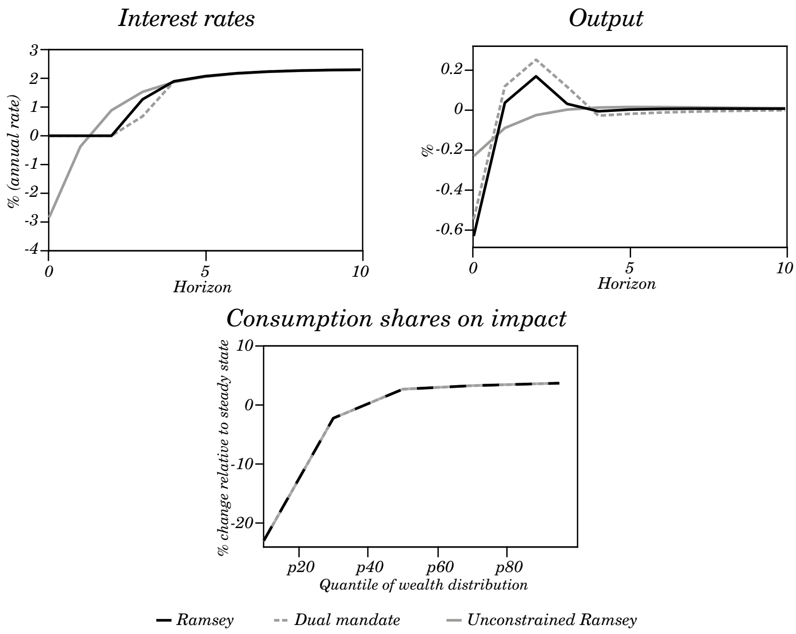

Source: Authors' calculations.

Note: Impulse responses of interest rates, output, and consumption shares (at $t = 0$) to the inequality shock under optimal monetary policy for the Ramsey planner subject to the ZLB (black), the dual-mandate policymaker subject to the ZLB (dashed grey), and the unconstrained Ramsey planner (solid grey).