Optimal Central Bank Balance Sheets with Incomplete Markets*

Zachary Bethune
Rice University

First version: May 2026; This version: September 2026
(Preliminary, please do not distribute.)

Abstract

I study the optimal size of a central bank's balance sheet in a continuous-time, incomplete-markets economy where households self-insure against income and expenditure risk by accumulating two nominal assets, liquid money and illiquid bonds. The size of the balance sheet determines the relative supply of the two assets in the hands of the public — since to create money the central bank must purchase bonds — and their real returns, given by the inflation rate and the liquidity premium. Optimal policy faces a trade-off between providing social insurance, financed by taxing wealth through low real returns, and incentivizing private self-insurance by engineering high real returns. The resolution depends on the ability of the central bank to commit. A central bank with commitment power runs a large real balance sheet that satiates the economy with liquidity, achieving a zero liquidity premium with mild long-run deflation to drive real returns high and reward self-insurance. A discretionary central bank instead restricts liquidity with a smaller real balance sheet, sustaining a positive liquidity premium and higher long-run inflation. I show this intuition survives extending the model to incorporate interest on money, productive real private capital, and richer portfolio adjustment costs and income process.

*Contact: bethune@rice.edu. I thank Marco Bassetto, Wei Cui, Ricardo Lagos, Guillaume Rocheteau, Venky Venkateswaran, Steve Williamson, and Russell Wong for helpful comments and suggestions. I also thank seminar and conference participants at the 2026 Summer Workshop on Money, Banking, Payments, and Finance in Gerzensee, LEMMA at Paris Pantheon Assas University, and Rice University.

1 Introduction

Since 2008, central bank balance sheets in the United States and other developed economies have expanded to historically high levels. Policymakers have increasingly called for reducing the size of these balance sheets, even though the process of balance sheet normalization has proven difficult in practice.[1] Much of the discussion has focused on a positive question about feasibility (Acharya et al., 2022; Duffie, 2026): can the central bank shrink its balance sheet while maintaining liquidity regulation and monetary policy implementation? This paper instead asks a normative question: should the central bank aim to have a smaller balance sheet?

I approach this question starting from a simple observation that the assets that comprise the two sides of a central bank's balance sheet differ in their characteristics. On the liability side, a central bank primarily issues short-term, money-like assets, such as reserves or currency in circulation. On the asset side, the composition includes a mix of government debt and private assets that tend to be less liquid, longer maturity, and more risky. By expanding or shrinking the balance sheet, the central bank influences the aggregate supply of assets on either side, and in turn, their prices and spreads. The goal of this paper is to study the determination of the optimal size of the balance sheet based on the distortions in asset prices and spreads it should support, which can be viewed as an optimal taxation problem in the spirit of Phelps (1973), Lucas and Stokey (1983), and Chari and Kehoe (1999).

This paper focuses on the liquidity difference between central bank assets and liabilities and the two prices they influence: the liquidity premium and the inflation rate. When the central bank expands its balance sheet, it buys less-liquid assets and pays for them by issuing money-like liabilities. The purchase makes liquid assets abundant relative to less-liquid ones, and the compensation the public requires for holding the less-liquid ones—the liquidity premium—falls.[2] The same transaction also moves the inflation rate. Since asset purchases are made with newly issued money, a path for the size of the balance sheet is a path for money creation, which becomes inflation unless fiscal policy absorbs it (Sargent and Wallace, 1981; Hall and Reis, 2015; Del Negro and Sims, 2015). Both the inflation and the liquidity premium are taxes on the assets the public holds. Inflation taxes nominal wealth while the liquidity premium is a tax on money.

I study the optimal balance sheet problem in a continuous-time, incomplete-markets economy in the tradition of Bewley (1980, 1983), with expenditure risk following Rocheteau et al. (2018a) and two assets as in Kaplan et al. (2018). Households face persistent shocks to labor productivity and random, idiosyncratic lumpy expenditure shocks. They self-insure by accumulating

[1] For instance, Federal Reserve Chair Kevin Warsh has stated that the Fed's balance sheet is "too big" and has called for a return to a smaller footprint (confirmation testimony, U.S. Senate Committee on Banking, Housing, and Urban Affairs, February 2025).

[2] For evidence that the liquidity premium moves with the supply of money-like assets, see Krishnamurthy and Vissing-Jorgensen (2012) and Nagel (2016). Copeland et al. (2025) document its return when the balance sheet shrank and repo rates spiked in September 2019.

two publicly-supplied nominal assets: money, which is liquid and can finance lumpy consumption, and bonds, which cannot. The liquidity premium compensates households for holding the illiquid asset, and the government's balance sheet determines the supply of both assets in the hands of the public.[3] I study optimal policy in two institutional arrangements: i) a consolidated government that can jointly set the supply of money and bonds as well as transfers versus ii) a monetary authority separated from the fiscal authority that issues money in exchange for bonds — equivalent to a restriction on remittances between the two authorities — and takes a fiscal rule as given. Within each arrangement, I solve for optimal policy under full commitment and under full discretion, following the approach of Roberds (1987) and Dávila and Schaab (2023).

To start, I consider the case in which fiscal transfers are rich, able to condition on each household's idiosyncratic state. In this case, the fiscal tax and transfer system can directly redistribute and provide insurance, except for lumpy expenditure risk subject to the liquidity constraint.[4] The optimal policy, under commitment or discretion, then satiates the economy with liquidity. The planner sets the relative bond supply low enough that the liquidity premium hits zero, and money earns the rate of time preference. This result is the classic Friedman (1969) rule which says, since public liquidity is socially costless to produce, it should be costless for private agents to hold.[5] The size of the balance sheet is indeterminate in the sense that the ratio of money to bonds in the hands of the public must be high enough to yield a zero liquidity premium, but setting it any larger is neutral. In other words, at the optimum, marginal open market operations are neutral, which echoes the irrelevance results of Wallace (1981) and Chamley and Polemarchakis (1984).

When transfers are restricted to lump sum, markets are incomplete and in any decentralized equilibrium marginal utilities of consumption are not equalized across households, they fall with wealth and income. This introduces a tradeoff in optimal policy between public and private insurance. Increasing transfers financed by inflation or low liquidity premia provides public insurance through redistribution. For instance, inflation taxes high nominal wealth households and redistributes through lump-sum transfers to low nominal wealth households. Reducing the liquidity premium, under a fixed inflation rate, taxes high bond-holding households and redistributes to lower bond-holding households, which in a decentralized equilibrium redistributes from those with low to high marginal utility. A planner who wants to provide more public insurance would be pushed toward higher inflation and a lower liquidity premium. This pol-

[3] Introducing lumpy expenditure risk in an otherwise frictionless two-asset economy is a way to introduce liquidity differences between assets while keeping the idiosyncratic state space small to only total wealth and income. In Section 6, I extend the optimal policy problem to an environment with continuous portfolio adjustment costs as in the Heterogeneous-Agent New Keynesian (HANK) literature following Kaplan et al. (2018) and show qualitatively the results carry over.

[4] This is true because I assume transfers are in flows while the liquidity constraint is on lumpy purchases.

[5] The optimality of the Friedman rule survives distorting taxes and redistributive motives when the tax system is rich enough (Kimbrough, 1986; Chari et al., 1996; da Costa and Werning, 2008).

icy, however, destroys the incentive for wealth accumulation and self-insurance since it reduces the real returns to savings. A planner that wants to encourage self-insurance would be pushed toward lower inflation and a higher liquidity premium.

How this tradeoff is resolved depends on the planner's ability to commit. I show, under discretion, the push for public insurance through redistribution wins and the planner increases inflation without bound. To a discretionary planner the inherited buffer stock of nominal wealth is sunk. Dilution through additional nominal asset creation looks like free redistribution since its cost is the wealth accumulation it deters through lower expected returns. This inflation temptation echoes the classic time-consistency problem (Calvo, 1978; Barro and Gordon, 1983), but its source here is a redistribution incentive under incomplete markets. Under commitment, the planner instead rewards self-insurance by pushing inflation close to the Friedman rule (deflation), stopping only when the lump-sum taxes required to finance it begin to reach their upper bound of the lowest income households. The commitment planner internalizes the wealth accumulation that the path of future policy induces. Higher sustained expected returns induce households to hold larger buffer stocks, substituting self-insurance for public insurance. The liquidity premium itself closes to zero, since deflation carries both the money and bond returns toward the rate of time preference.

I then turn to studying optimal policy when the central bank is separated from the Treasury. The central bank controls the money supply through open market operations—swapping money for bonds on its balance sheet—but makes no direct transfers. Fiscal policy sets transfers and debt issuance by a debt-stabilizing rule, as in Leeper (1991). Separation changes the nature of the optimal policy problem by removing the independence of the two tax instruments: inflation and the liquidity premium. Whereas the consolidated government can independently determine inflation and liquidity premium, the separated central bank's budget constraint ties them together. Higher inflation lowers real money balances and with them the real size of the balance sheet, leaving more real bonds in the hands of the public and a higher liquidity premium. Hence, a planner that wants high inflation must also accept a higher liquidity premium.

Under discretion, the separated central bank faces a different redistributive calculation compared to the consolidated government. Since higher inflation also raises the liquidity premium, the inflation tax falls disproportionately on money holdings, which are less dispersed than total wealth. Inflation has less redistributive power and the discretionary optimum of the separated central bank stops inflating when the smaller redistributive gain balances the cost of more expensive liquidity.[6] The optimal liquidity premium is strictly positive, which implies the balance sheet is set so that liquidity is not abundant.

[6]The result that separating the monetary and fiscal authority leads to a smaller discretionary incentive for inflation echoes Sargent (1982) in which the hyperinflations in Austria, Hungary, Poland, and Germany each ended abruptly when the government paired a change in the fiscal regime with an independent central bank legally barred from financing the government. In the model, that pair of measures requires no commitment power.

Under commitment, the separated central bank is also limited but, as in the consolidated problem, is pushed the other way from the discretionary planner and aims to incentivize self-insurance. It optimally satiates the economy with liquidity by setting inflation low enough so that the liquidity premium reaches zero. Lowering inflation raises real money demand, grows the real balance sheet, and pulls the premium down.

I calibrate the model to match features of the wealth and income distribution in the United States. Income risk is set to match flows into and out of unemployment and the replacement rate of unemployment insurance. Expenditure risk is set to match the level and interest sensitivity of money demand, and the supplies of money and bonds match the ratio of government liabilities held by domestic private investors to GDP of fifty-two percent and a liquidity premium of two percent. Under discretion, the separated central bank's stationary optimum inflation is 1.79 percent and the optimal liquidity premium is 4.17 percent. Liquidity remains scarce at the optimum. Under commitment, the central bank reaches the zero-premium corner at a deflation rate of half a percent. In Section 5, I study the optimal balance sheet along transitions following aggregate shocks, a productivity contraction and a rise in unemployment risk. In both cases it is optimal to keep the size of the balance sheet relatively fixed and let the two prices absorb the shock. The adjustment is spread across inflation and the liquidity premium, cheapening liquidity when households value it most. Fixing either price instead relocates the burden of the shock onto the other margin, forces large movements in the size of the balance sheet, and raises consumption dispersion. The optimal balance sheet absorbs shocks through its prices rather than its size.

Section 6 examines the robustness to three extensions. First, I show allowing the central bank to pay interest on money, as the Federal Reserve has since 2008, does not alter the policy prescriptions. Any pair of a nominal rate on money and an inflation rate that delivers the same real return as in the baseline implements the same optimum. Additionally, the promised nominal interest rate on money is still a nominal claim that dilution reaches, so the temptation to inflate is unchanged. Second, I add productive capital which anchors the return on illiquid assets to the marginal product of capital, so inflation passes into the liquidity premium one for one so even a consolidated planner faces the trade-off that only separation created in the baseline. The asymmetry between discretion and commitment survives even though capital shrinks the nominal tax base by an order of magnitude. Third, I consider portfolio adjustment costs and a richer income process which jointly match the empirical wealth distribution. Optimal policy keeps its qualitative features. Discretion still leans toward inflationary redistribution while commitment still leans toward rewarding self-insurance.

The rest of the paper proceeds as follows. Section 1.1 summarizes the contribution relative to the literature. Section 2 presents the model and characterizes the decentralized equilibrium. Section 3 studies optimal policy for the consolidated government, under discretion and under commitment. Section 4 studies the optimal balance sheet of a central bank separated from the

Treasury. Section 5 studies the optimal response of the separated central bank to aggregate shocks along transitions. Section 6 presents the three robustness exercises, and Section 7 concludes.

1.1 Literature

A long literature studies the optimal supply of government debt in incomplete-markets economies, where public liabilities provide liquidity, self-insurance, and safe collateral that private markets cannot fully replicate (Woodford, 1990; Aiyagari and McGrattan, 1998; Holmström and Tirole, 1998; Azzimonti and Yared, 2019; Gorton and Ordoñez, 2022; Angeletos et al., 2023; Aguiar et al., 2024). Bhandari et al. (2017a,b) extend the question to the dynamic management of the government's debt portfolio, chosen to hedge aggregate fiscal risk. That literature tends to treat public liabilities as a single asset, or as a portfolio distinguished by returns, while this paper starts from the observation that central bank liabilities (money) and Treasury debt (bonds) differ in their liquidity properties, so the optimal-supply question is jointly one of quantity and composition. In this way, the paper relates to another long-standing literature on the optimum quantity of money, in which money is the vehicle for self-insurance (Bewley, 1980, 1983; Kehoe et al., 1990; Molico, 2006; Bhattacharya et al., 2005; da Costa and Werning, 2008; Lippi et al., 2015; Chien and Wen, 2024). Relative to both, I characterize the joint optimum over the quantity and the composition of the government's nominal liabilities—the optimal inflation rate and the optimal liquidity premium—under both commitment and discretion.

The asset structure follows the New Monetarist literature on the coexistence of money and less-liquid assets (Lagos and Wright, 2005; Rocheteau and Wright, 2005; Lagos et al., 2017), with lumpy expenditure risk within an incomplete-markets framework as in Rocheteau et al. (2018a), who study money alongside illiquid bonds and show that a positive liquidity premium and a positive inflation rate can maximize steady-state welfare, but do not solve a full optimal-policy problem. The result that a positive liquidity premium can be socially optimal connects the paper to the literature on the societal role of illiquid bonds, including Kocherlakota (2003), Boel and Camera (2006), Berentsen et al. (2007), Shi (2008), and Andolfatto (2011). By assuming bonds and money differ in their liquidity, I connect to the literature on open market operations. Rocheteau et al. (2018b) show that liquidity differences give open market operations real effects, and Sterk and Tenreyro (2018) and Bustamante (2023) study their redistributive incidence through the revaluation of nominal government debt or through households' portfolios of money and bonds. Here those channels become the objects of optimal policy design and the balance sheet is treated as a permanent policy instrument. Bianchi and Bigio (2022) and Bigio and Sannikov (2026) study balance-sheet policy and its implementation through the banking sector. In Bianchi and Bigio (2022), banks fund illiquid loans with deposits and hold precautionary reserves against withdrawal risk, so the central bank's balance sheet moves lending by repricing that liquidity risk.

Bigio and Sannikov (2026) embed this machinery in an incomplete-markets economy and show that control of the credit spread carries a trade-off between insurance across households and insurance against aggregate shocks. They pose solving for the optimal balance sheet as an open question. This paper takes up that question in a complementary setting. The liquidity margin here sits in household portfolios rather than in bank settlement, and the precautionary demand the balance sheet prices is the household's own.[7]

Methodologically the paper follows the continuous-time heterogeneous-agent optimal-policy literature (Achdou et al., 2022; Nuño and Thomas, 2022; Dávila and Schaab, 2023; Bhandari et al., 2026; Le Grand et al., 2022; Challe, 2020; Acharya and Dogra, 2020; Hurtado et al., 2023), adopting the approach of Dávila and Schaab (2023). I depart by working in a Bewley-type model where money is essential and the instruments are the nominal quantities on the government balance sheet, activating seigniorage and revaluation channels.[8] Closest on the revaluation margin, Bhandari et al. (2021) study surprise revaluations of household nominal claims as state-contingent fiscal insurance under commitment, whereas here the same margin is a permanent redistributive levy whose incidence the optimality conditions price and whose time consistency is the central question.

The discretionary analysis related to the literature on time-consistent monetary and fiscal policy, from the inflation bias of Kydland and Prescott (1977), Calvo (1978), and Barro and Gordon (1983)—here redistributive rather than expansionary—to the time consistency of fiscal–monetary plans and their Markov-perfect equilibria (Lucas and Stokey, 1983; Obstfeld, 1997; Chang, 1998; Albanesi et al., 2003; Persson et al., 2006; Alvarez et al., 2004; Klein et al., 2008; Niemann et al., 2013; Straub and Werning, 2020). Within this tradition, the closest work gives public liabilities distinct liquidity roles. Martin (2009, 2011, 2013, 2015) studies Markov-perfect policy in monetary economies where money is essential in exchange and government debt is not, so the composition of public liabilities disciplines the discretionary equilibrium. Bassetto and Cui (2024) make the liquidity services of public debt the centerpiece of a Ramsey plan. There the liquidity structure disciplines the government's financing problem. Here it prices household self-insurance, and the redistributive temptation it creates across the wealth distribution is the time-consistency problem. In Díaz-Giménez et al. (2008) nominal government debt is the burden tempting a Markov-perfect authority, and Hurtado et al. (2023) study that temptation in continuous time, with inflation eroding the government's own long-term debt. Nuño and Thomas (2022) is the closest treatment of redistributive dilution, where discretionary inflation redistributes across a household distri-

[7] A related literature studies monetary policy implementation and the central bank's operating framework, where balance-sheet policy transmits through banks and asset markets rather than through its incidence across the household wealth distribution. For instance, see Martin and Monnet (2011), Lagos and Navarro (2023), Benigno (2020), Benigno and Nisticò (2020), Herrenbrueck (2019), Schulhofer-Wohl (2025), and Williamson (2025); on floor versus corridor operating systems, Cúrdia and Woodford (2011) and Arce et al. (2020); on asset purchases, Gertler and Karadi (2011) and Greenwood et al. (2015).

[8] Kaplan (2025) catalogs these fiscal–monetary interconnections in HANK models.

bution of nominal positions. Here the diluted claims are the public liabilities themselves, they differ in liquidity, and the temptation's measure is a covariance. Separating monetary from fiscal policy connects to the delegation tradition (Rogoff, 1985; Walsh, 1995), to work on central-bank solvency, remittances, and fiscal backing (Wallace, 1981; Del Negro and Sims, 2015; Hall and Reis, 2015; Benigno and Nisticò, 2017), and to the fiscal–monetary regime literature (Sargent and Wallace, 1981; Leeper, 1991; Sims, 1994; Woodford, 1995; Bassetto, 2002).[9]

2 Model

I consider a continuous-time, incomplete-markets economy with two publicly-supplied nominal assets: money and bonds. Households receive idiosyncratic labor-productivity and lumpy-expenditure shocks, and self-insure by accumulating wealth in money and bonds. Money is liquid in the sense that it is used as a transaction account and means of payment for consumption. There exist continuously-open, frictionless markets in which households can trade money for bonds and, likewise, money for a *flow* of consumption goods. However, in the event of a lumpy expenditure shock, I assume that households can only access money (bonds are illiquid).[10] In equilibrium, bonds earn a weakly higher return than money—a liquidity premium that compensates households for the inability to use bonds in all transactions. Policy takes the form of a portfolio of public, nominal assets and a transfer rule.

2.1 Households

Time is continuous, $t \in [0, \infty)$, and there is a single, perishable consumption good. There exists a unit measure of households that discount the future at rate $\rho > 0$ and value flows of consumption, c_t, according to the instantaneous utility function $u(c_t)$. In addition to consuming in flows, households receive random, idiosyncratic preference shocks to consume discrete (lumpy) amounts of the consumption good, denoted y_t. These shocks arrive at Poisson rate $\alpha > 0$ and consuming y_t delivers the household a lump of utility $U(y_t)$.

Given adapted, left-continuous processes (c_t, y_t), a household's lifetime expected discounted

[9] A recent literature re-examines price-level determination with non-Ricardian (heterogeneous) households: Hagedorn (2017) develops a demand theory of the price level; Hagedorn (2024b,a) argues that the fiscal theory's foundations are fragile; and Angeletos et al. (2024) show that deficit-driven inflation in HANK economies operates through a classical non-Ricardian demand channel rather than the fiscal theory's valuation equation. The nominal anchor used here—quantity targeting combined with real fiscal transfers (Section 2.3)—sits in this non-Ricardian determinacy tradition.

[10] An alternative formulation is to have portfolio rebalancing as a continuous flow d_t subject to a strictly convex adjustment cost $\Phi(d_t)$, following Kaplan and Violante (2014) and Kaplan et al. (2018). I present this alternative in Appendix G and show that the optimal policy results are similar.

utility at time zero is given by

$$\mathbb{E}\left[\int_0^\infty e^{-\rho t}\, u(c_t)\, dt + \sum_{n=1}^{\infty} e^{-\rho T_n} U(y_{T_n})\right], \tag{1}$$

where the sequence $\{T_n\}_{n=1}^{\infty}$ represents the Poisson arrival times of lumpy expenditure shocks. The functions u and U are strictly increasing and concave, and satisfy the Inada conditions $u'(0) = U'(0) = \infty$ and $u'(\infty) = U'(\infty) = 0$. Households supply one unit of labor inelastically and face idiosyncratic labor productivity given by an exogenous Markov process z_t. Let $\mathcal{Z}$ stand for the support of z. A household's labor earnings are $z_t w_t$, where w_t is the wage rate in a competitive labor market. There is a representative firm that produces output Y_t using a linear production technology $Y_t = L_t$, where L_t is aggregate labor.

The first-best allocation under a utilitarian social welfare function is

$$u'(c^{FB}) = U'(y^{FB}), \tag{2}$$

$$c^{FB} + \alpha\, y^{FB} = \mathbb{E}_z[z]\,, \tag{3}$$

where $\mathbb{E}_z$ denotes an expectation under the ergodic distribution of productivity. In (2), the planner equalizes the marginal utility of consumption across agents, states, and dates, using total resources, subject to the resource constraint (3).

2.2 Assets and government

There are two publicly-supplied, risk-free nominal assets: money M_t and bonds B_t. Both are nominal claims on the same unit of account and trade at real price ϕ_t per nominal unit. Bonds pay a nominal interest rate $i_t^b \geq 0$, whereas money pays no interest.[11] The corresponding real returns are

$$r_t^a = \dot{\phi}_t/\phi_t \equiv -\pi_t, \qquad r_t^b = i_t^b + \dot{\phi}_t/\phi_t,$$

where π_t is the rate of inflation and $i_t^b = r_t^b - r_t^a \geq 0$ is the liquidity premium.

From the standpoint of households, the assets differ in only one substantive way: money can be used to finance lumpy expenditures and bonds cannot. At every instant between Poisson lumpy shocks, the household freely allocates total real wealth between money and bonds at zero cost. At the moment a lumpy shock arrives, however, only money already held can be used to

[11] I assume the government is not allowed to pay interest on money following the classical literature in monetary theory, e.g. Lucas and Stokey (1983). One can view money as central bank liabilities (currency in circulation and reserves) paired with a competitive narrow banking system that converts these one for one into deposits or other transaction accounts that households hold. I study the possibility of allowing interest on money (in line with the practice of paying interest on reserves) in Section 6.1.

finance the purchase. Let ω_t denote a household's total real wealth, a_t its real money holdings, and $b_t = \omega_t - a_t$ its real bond holdings. I assume that money and bond holdings are non-negative, $a_t, b_t \geq 0$. The household's law of motion for total real wealth (outside lumpy-expenditure events) is

$$\dot{\omega}_t \;=\; r_t^a a_t + r_t^b b_t + z_t - c_t + \mathcal{T}_t \;=\; r_t^b \omega_t - i_t^b a_t + z_t - c_t + \mathcal{T}_t, \tag{4}$$

where the second equality uses $r_t^b = r_t^a + i_t^b$. Holding wealth in money rather than bonds sacrifices the bond return at rate i_t^b. At the moment of a lumpy expenditure shock, the household consumes $y_t \leq a_t$, and total wealth jumps to $\omega_t - y_t$.[12]

I distinguish between the policy of the fiscal authority (Treasury) and that of the monetary authority (central bank). Let B_t^T denote the total outstanding nominal debt of the Treasury, which is held either by the public, B_t, or by the central bank, B_t^{CB}, so that $B_t^T = B_t + B_t^{CB}$. The Treasury sets transfers $\mathcal{T}_t(\omega, z)$, that can generically depend on a household's state, and issues debt B_t^T. Given a remittance R_t from the central bank, the Treasury's budget constraint is

$$\int \mathcal{T}_t(\omega, z)\, dG_t + i_t^b \phi_t B_t^T = R_t + \phi_t \dot{B}_t^T. \tag{5}$$

Since the household freely allocates wealth between money and bonds at every instant, the portfolio split (a, b) is endogenously pinned down by (ω, z) in equilibrium, so there is no loss of generality in conditioning transfers on total wealth rather than on the portfolio composition. The left-hand side of (5) represents the Treasury's total expenditures from transfers and interest on Treasury debt, while the right-hand side represents its total revenue from central bank remittances and debt issuance.

The central bank sets the supply of money M_t by conducting open market operations—swapping money for bonds on its balance sheet. Its budget constraint is

$$R_t + \phi_t \dot{B}_t^{CB} = i_t^b \phi_t B_t^{CB} + \phi_t \dot{M}_t. \tag{6}$$

The left-hand side of (6) represents the central bank's expenditures from remittances and bond purchases, while the right-hand side represents its total revenue from money creation and interest on Treasury debt.

Adding (5) and (6), the remittance R_t cancels and the *consolidated government's* budget con-

[12] To make money essential for trade, one can imagine that households can commit to deliver their bond positions across time but cannot access their commitment technology in the instant they receive a lumpy expenditure shock. Further, a long position of bonds cannot be authenticated quickly. See Rocheteau et al. (2018a) for a related discussion.

straint is

$$\int \mathcal{T}_t(\omega, z)\, dG_t + i_t^b \phi_t B_t = \phi_t \dot{M}_t + \phi_t \dot{B}_t. \tag{7}$$

When monetary and fiscal policy are jointly optimized (Section 3), the relevant constraint is (7). When they are separated (Section 4), the individual constraints (5)–(6) are considered separately along with a fiscal policy rule.

In addition to the flow budget constraint, the government is subject to a *no-Ponzi condition* ruling out indefinite accumulation of real liabilities:

$$\lim_{t\to\infty} e^{-\rho t}\phi_t(M_t + B_t) \leq 0. \tag{8}$$

Combined with (7) and the household budget constraints aggregated over the distribution of households, condition (8) delivers the intertemporal government budget constraint: the present value of real government liabilities equals the present value of future primary surpluses (transfers, net of seigniorage and interest costs).[13]

2.3 Decentralized equilibrium

Given a linear production function, the labor market clearing condition requires $w_t = 1$ for all t. Households take policy and prices (ϕ_t, i_t^b) as given. Let $W_t(\omega, z)$ represent a household's value function at time t, with ω total real wealth and z idiosyncratic productivity. It satisfies the Hamilton-Jacobi-Bellman (HJB) equation

$$\rho W_t(\omega, z) = \max_{c_t, y_t, a_t} \; u(c_t) + \alpha U(y_t) + \tilde{\mathcal{A}}_t(\omega, z, c, y, a)\,[W_t] + \partial_t W_t, \tag{9}$$

where the constraints on choices are $c_t \geq 0$, $0 \leq y_t \leq a_t$, $a_t \in [0, \omega_t]$ (equivalently $b_t = \omega_t - a_t \geq 0$), and the state constraint $\omega_t \geq 0$. The operator $\tilde{\mathcal{A}}_t$ represents the infinitesimal generator of the idiosyncratic state (ω, z) and is given by

$$\tilde{\mathcal{A}}_t[W] = \left[r_t^b \omega - i_t^b a_t + z_t - c + \mathcal{T}_t\right] \partial_\omega W + \alpha\,[W(\omega - y, z) - W(\omega, z)] + \mathcal{A}_z W(\omega, z). \tag{10}$$

The generator has three components: continuous changes in total real wealth (first term), discrete jumps from lumpy consumption at a Poisson rate (second term), and idiosyncratic productivity transitions (third term).

The household has three controls: flow consumption c_t, intended lumpy consumption y_t, and

[13] Household optimization combined with the debt limit $b_t \geq 0$ implies the standard household transversality $\lim_{t\to\infty} e^{-\rho t}\partial_\omega W_t \cdot (\omega_t) = 0$. Together with (8), this rules out paths where real wealth or real liabilities grow at a rate exceeding ρ.

money holdings a_t. The optimization problem leads to three optimality conditions. The first is the flow consumption Euler:

$$u'(c_t) = \partial_\omega W_t(\omega, z). \tag{11}$$

The second governs the money-holdings choice. Differentiating with respect to a_t in the HJB and applying complementary slackness on the bounds $y_t \leq a_t$ and $a_t \leq \omega_t$ yields

$$i_t^b \, \partial_\omega W_t(\omega, z) \;=\; \xi_t - \zeta_t, \tag{12}$$

where $\xi_t \geq 0$ is the multiplier on $y_t \leq a_t$ (with $\xi_t(a_t - y_t) = 0$) and $\zeta_t \geq 0$ is the multiplier on $a_t \leq \omega_t$ (with $\zeta_t(\omega_t - a_t) = 0$). When the liquidity premium is strictly positive ($i_t^b > 0$), the LHS is strictly positive since $\partial_\omega W_t > 0$, and so $\xi_t > \zeta_t \geq 0$, which forces $a_t = y_t$: the household holds money exactly equal to its intended lumpy purchase. Holding money in excess of y_t is strictly costly (the household marginally sacrifices $i_t^b \, \partial_\omega W_t$) and provides no marginal benefit, because the household will not consume more than its planned y_t at the shock. When $i_t^b = 0$ (the liquidity-trap case), the constraint $y_t \leq a_t$ need not bind and the household is indifferent over any $a_t \in [y_t, \omega_t]$. Without loss of generality, let $a_t = y_t$ in this case also.

The third optimality condition is the first-order condition for lumpy consumption. Differentiating with respect to y_t and substituting the binding $a_t = y_t$:

$$\alpha[U'(y_t) - \partial_\omega W_t(\omega - y_t, z)] \;\geq\; i_t^b \, \partial_\omega W_t(\omega, z), \tag{13}$$

with equality if $y_t < \omega_t$. The LHS is the marginal expected lumpy benefit per unit time—the α-weighted gap between the marginal utility of the lumpy purchase $U'(y_t)$ and the post-shock marginal value of wealth $\partial_\omega W_t(\omega - y_t, z)$. The RHS is the marginal cost of holding wealth in money (rather than bonds), $i_t^b \, \partial_\omega W_t(\omega, z)$. When $y_t = \omega_t$ (full depletion), the upper bound on money holdings $a_t \leq \omega_t$ binds, the household holds no bonds, and the FOC holds as a strict inequality.

Substituting the household's optimal $a_t = y_t$ into the wealth dynamics, the law of motion simplifies to $\dot{\omega}_t = r_t^b \omega_t - i_t^b y_t + z_t - c_t + \mathcal{T}_t$. With the optimal policies substituted, the HJB takes the compact form

$$\rho W_t = u(c_t) + \alpha U(y_t) + \mathcal{A}_t \, [W_t] + \partial_t W_t, \tag{14}$$

where $\mathcal{A}_t = \tilde{\mathcal{A}}_t(\omega, z, c_t, y_t, y_t)$. The equilibrium measure G_t over household states (ω, z) satisfies

the Kolmogorov forward equation (KFE):

$$\partial_t G_t = \mathcal{A}_t^*[G_t], \tag{15}$$

where $\mathcal{A}_t^*$ is the adjoint of the generator $\mathcal{A}_t$. The stationary distribution may contain a mass point at the borrowing constraint $\omega = 0$ for each z. The mass point is denoted by $m_t(z)$, and the absolutely continuous part of G_t is given by the density $g_t(\omega, z)$. Let $F_t(\tilde{\omega}, z)$ stand for the measure over nominal wealth $\tilde{\omega}_t = \omega_t/\phi_t$ and productivity.

Market clearing requires real money demand to equal real money supply, and likewise for bonds:

$$\int y_t(\omega, z)\, dG_t(\omega, z) = \phi_t M_t, \tag{16}$$

$$\int [\omega - y_t(\omega, z)]\; dG_t(\omega, z) = \phi_t B_t, \tag{17}$$

where $a_t(\omega, z) = y_t(\omega, z)$ and $b_t(\omega, z) = \omega - y_t(\omega, z)$. Collect government policy starting in time t as $\mathbf{P}_t = \{\dot{M}_s, \dot{B}_s, \mathcal{T}_s(\omega, z)\}_{s \geq t}$. The decentralized equilibrium is defined as follows.

Definition 1 (Decentralized competitive equilibrium)**.** *Taking as given an initial nominal wealth distribution $F_0(\tilde{\omega}, z)$, initial nominal asset supplies $M_0 > 0$ and $B_0 > 0$, and a path of policy $\mathbf{P}_0$, a decentralized competitive equilibrium consists of value functions and decision rules $\{W_t, c_t, y_t\}_{t \geq 0}$, equilibrium measures $\{G_t\}_{t \geq 0}$, and prices $\{\phi_t, i_t^b\}_{t \geq 0}$, such that:*

i) value functions and decision rules satisfy (11)–(14) with $a_t = y_t$ and $b_t = \omega_t - y_t$,

ii) the equilibrium measure G_t satisfies (15),

iii) the government budget constraint (7) holds at each date t,

iv) and both market clearing conditions (16)–(17) hold at each date t.

Equations (7) and (11)–(17) serve as implementability constraints for optimal policy. Market clearing for labor is implemented by setting $w_t = 1$, and the market clearing condition for goods is satisfied by Walras' law when both asset market clearings (16)–(17) and the government budget constraint (7) hold, with total output given by $Y_t = L_t \equiv \int z_t\, dG_t(\omega, z)$.

Determination of the initial price level ϕ_0. A subtle but important feature of Definition 1 is how the initial price level ϕ_0 is determined by the equilibrium conditions. First, note that aggregate wealth market clearing at $t = 0$ reduces to the tautology $\phi_0(M_0 + B_0) = \phi_0 \int \tilde{\omega}\, dF_0$, given consistency of the inherited nominal distribution F_0 with the total nominal asset stock $M_0 + B_0$. The substantive equilibrium conditions at $t = 0$ are the money market clearing condition (16)

together with the government budget constraint (7):

$$\int y_0(\omega, z; i_0^b)\, dG_0 \;=\; \phi_0 M_0, \qquad \int \mathcal{T}_0(x)\, dG_0 + i_0^b\, \phi_0 B_0 = \phi_0 \dot{M}_0 + \phi_0 \dot{B}_0. \tag{18}$$

These two equations jointly determine the pair (ϕ_0, i_0^b). With $\mathcal{T}_0(x), \dot{M}_0, \dot{B}_0, M_0, B_0$ all specified as policy or initial conditions (and $\mathcal{T}_t$ in real terms), this joint system has a unique solution among the non-explosive equilibria.[14] The price path $\{\phi_t, i_t^b\}_{t>0}$ is then determined by market clearing at each subsequent date.[15]

2.4 Characterization of the decentralized equilibrium

Before studying optimal policy, it is useful to first establish some properties of the decentralized equilibrium. I impose the following condition on the transfer rule that guarantees all households can finance a strictly positive consumption flow, a condition that will be slack in the optimal policy problem.

Assumption 1 (Feasibility)**.** *For all t, transfers and the lowest productivity state satisfy $\underline{z}_t + \mathcal{T}_t(0, \underline{z}) > 0$.*

Assumption 1 rules out confiscatory transfers that would leave a household at the constraint $\omega = 0$ (i.e., $a = b = 0$) with non-positive resources. The characterization of the equilibrium begins with establishing that the value function is a constrained viscosity solution of the HJB equation on $[0, \infty) \times \mathcal{Z}$, deriving the boundary condition at $\omega = 0$, and establishing some useful properties of the value function W_t. Proofs of all propositions and lemmas in the paper are collected in Appendix A.

Lemma 1 (Value function properties and boundary conditions)**.** *Under Assumption 1:*

i) $W_t(\omega, z)$ is a constrained viscosity solution of the HJB equation (14) on $[0, \infty) \times \mathcal{Z}$.

ii) Boundary condition. For each t and $z \in \mathcal{Z}$, $\partial_\omega W_t(0, z) < \infty$. At $\omega = 0$, the consumption Euler $u'(c_t) = \partial_\omega W_t$ holds and the state-constraint boundary condition satisfies

$$\partial_\omega W_t(0, z) \geq u'\left(z_t + \mathcal{T}_t(0, z)\right), \tag{19}$$

with equality iff $\dot{\omega}_t(0, z) = 0$.

[14] I restrict attention to equilibria that converge to a steady state with finite inflation. Hyperinflationary paths with $\phi_t \to 0$ and a non-monetary equilibrium also exist under the same policy path, as in any economy with unbacked nominal assets. Only under commitment can the selection be supported by policy, through an arbitrarily small promise to redeem money at a price floor (Obstfeld and Rogoff, 1983).

[15] The fact that real transfers $\mathcal{T}_t$ are exogenously specified policy — from the standpoint of the decentralized equilibrium — is essential for this argument: if instead $\mathcal{T}_t$ adjusted passively to balance the budget at any price level (a "Ricardian" regime in the language of Leeper, 1991 and Woodford, 1995), ϕ_0 would be indeterminate. Conversely, if the central bank instead pegged a price (i_t^b, say) and let nominal quantities adjust passively, classical indeterminacy results (Sargent and Wallace, 1975) would arise. Quantity targeting combined with real fiscal transfers leads to a nominal anchor.

iii) Monotonicity and concavity: W_t *is strictly increasing and strictly concave in* ω *for each* z.

iv) Asymptotic: $\lim_{\omega\to\infty} \partial_\omega W_t(\omega, z) = 0$.

Lemma 1 extends Rocheteau et al. (2018a, Theorem 1) to a non-stationary environment with idiosyncratic productivity risk. Part (ii) shows that households at the no-borrowing constraint have a strictly positive but bounded marginal value of additional wealth. Part (iii) establishes the strict monotonicity and concavity of the value function, which will deliver a redistributive motive for optimal policy. Part (iv) is the standard transversality property in ω. The next result also extends Rocheteau et al. (2018a, Proposition 1).

Lemma 2 (Endogenous segmentation)**.** *For each* $t \geq 0$ *and* $z \in \mathcal{Z}$*, there exists a threshold* $\underline{\omega}_t(z) > 0$ *such that the household optimally chooses* $y_t(\omega, z) = \omega$ *for all* $\omega \in (0, \underline{\omega}_t(z)]$ *(holds all wealth in money), where the threshold satisfies*

$$\alpha\left[U'(\underline{\omega}_t(z)) - \partial_\omega W_t(0, z)\right] \;=\; i_t^b\, \partial_\omega W_t(\underline{\omega}_t(z), z). \tag{20}$$

For $\omega > \underline{\omega}_t(z)$*,* $y_t(\omega, z) < \omega$ *and the lumpy FOC* (13) *holds with equality. For each* $\omega > 0$*,* $0 \leq \partial_\omega y_t(\omega, z) \leq 1$.

Below the threshold $\underline{\omega}_t(z)$, the household is too wealth-poor to keep any positive bond holdings. It holds all of its wealth in money, $a_t = \omega_t$. In the event of a lumpy shock, it fully depletes its money balance ($y_t = a_t = \omega_t$), leaving total wealth at the borrowing constraint $\omega = 0$. The lumpy FOC (13) holds with slack in this regime because the upper-bound constraint $a_t \leq \omega_t$ binds and contributes a positive Kuhn-Tucker multiplier. Above the threshold, the household holds both money $a_t = y_t < \omega_t$ and bonds $b_t = \omega_t - y_t > 0$, the lumpy FOC (13) holds with strict equality, and the household partially depletes money while keeping a positive bond position. This endogenous segmentation—wealth-poor households hold all wealth in money, wealth-rich households hold both money and positive bonds—is a key feature of the two-asset structure relative to the pure-currency benchmark.

Stationary equilibrium. This section discusses the existence and uniqueness of a stationary equilibrium given a stationary policy, with the restriction to lump-sum transfers.[16] A policy is stationary if both nominal stocks grow at a common rate g^M and real lump-sum transfers are constant. A stationary equilibrium under a stationary policy is then a decentralized equilibrium in which all policies and value functions, the distribution of real household wealth and income, $G_t(\omega, z)$, and real returns, r_t^b and r_t^a (and hence i_t^b), are constant. Constancy of G_t and market clearing imply that real balances $\phi_t M_t$ and real bonds $\phi_t B_t$ are constant, so the growth rate

[16] Most of the optimal policy analysis restricts transfers to be lump-sum. For existence and uniqueness, any transfer schedule satisfying Assumption 1 suffices.

determines inflation, $g^M = -\dot{\phi}_t/\phi_t = \pi$. The common growth rate across nominal assets implies the composition $\beta = B_t/M_t$ is constant at its initial level, chosen by policy. A stationary policy can then be described by a pair (g^M, β).

Lemma 3 (Existence of stationary equilibrium)**.** *Fix a stationary policy (g^M, β) with $g^M > -\rho$ and $\beta \geq 0$, and let the constant transfer implied by the budget satisfy Assumption 1. There exist thresholds $0 < \underline{\beta}(g^M) \leq \bar{\beta}(g^M) \leq \infty$ such that a stationary monetary equilibrium exists for every $\beta \in [0, \bar{\beta}(g^M))$. The equilibrium premium is strictly positive if $\beta \in (\underline{\beta}(g^M), \bar{\beta}(g^M))$ and zero if $\beta \leq \underline{\beta}(g^M)$.*

Rocheteau et al. (2018a) establish existence in the same two-asset economy without idiosyncratic income risk, including the threshold structure of the zero-premium region. At fixed inflation, the feasible premia lie between zero and $\rho + g^M$. At the upper end, the real bond return approaches the rate of time preference and the demanded ratio of bonds to money rises toward its upper limit $\bar{\beta}(g^M)$, which may be infinite. Uniqueness, left unproven, holds if the demanded bond–money ratio rises strictly with the premium. In partial equilibrium, a higher premium raises the return on bonds and household portfolios shift toward them. The difficulty in proving uniqueness is that, in general equilibrium, the transfer adjusts with the liquidity premium and feeds back into demand. Uniqueness requires that the wealth demand not fall too strongly in the lump-sum transfer. The quantitative results explicitly check for multiplicity of this kind and find none.

2.5 Comparative statics: π versus i^b

Finally, before introducing the optimal policy problem, it is instructive to study how stationary equilibrium allocations vary with changes in the inflation rate π and the liquidity premium i^b (induced by changes in government policy (g^M, β)). Specifically, Figures 1 and 2 trace the stationary equilibrium along each instrument, holding the other fixed and letting lump-sum transfers adjust to balance the government budget constraint. First, I introduce a calibration that will be used throughout the paper.

Calibration Normalize a unit of time to one year. Flow and lumpy utility are CRRA with curvature parameters $\sigma = 2$ and σ_y, respectively. Income follows a two-state employment process. Employed households earn $z_e = 1$ and unemployed households earn $z_u = 0.5$, targeting a fifty percent replacement rate. The job-finding rate $\lambda_f = 4.28$ per year matches a monthly job-finding probability of 30%, and the separation rate of $\lambda_s = 0.27$ delivers a steady-state unemployment rate of six percent, in line with postwar averages (Shimer, 2005).

Money in the model is treated as transaction balances and its empirical counterpart as the adjusted M1 aggregate of Lucas and Nicolini (2015)—M1 plus money-market deposit accounts—whose demand is stable over the postwar period. Money growth is set at $g^M = 2\%$ so stationary

Parameter		Value	Target / source
Independently calibrated			
σ	risk aversion	2	standard
Σ_y	lumpy-utility scale	1	normalization
z_e, z_u	income states	1, 0.5	fifty percent replacement ratio
λ_f, λ_s	job-finding, separation	4.28, 0.27	monthly job-finding 0.30; $u = 6\%$
g^M	money growth	2%	two percent inflation
Jointly calibrated			
α	lumpy-expenditure arrival	0.49	$\}$ $Q/Y = 0.30$; semi-elasticity = 2.1
σ_y	lumpy-utility curvature	0.45	
ρ	discount rate	3.4%	Govt liabilities to GDP $\phi(M+B)/Y = 0.52$
B/M	balance-sheet composition	0.73	$i^b = 2\%$

Table 1: Baseline calibration.

inflation is two percent. Bonds are treated as the stock of marketable Treasury debt held by domestic private investors, including both short-term and long-term debt.[17]

The remaining parameters are jointly calibrated. The lumpy-expenditure parameters (α, σ_y) target the level and interest sensitivity of money demand. Transaction balances averaged thirty percent of GDP post-1980, with an interest semi-elasticity of 2.1 (Ireland, 2009). This results in $\alpha = 0.49$ and $\sigma_y = 0.45$, which imply an expenditure shock roughly every two years with an average size of thirty percent of annual income—relatively large but infrequent expenditure risk. The discount rate ρ is set to target government liabilities held by domestic private investors at fifty-two percent of GDP. Money is held domestically by construction. For bonds, I exclude the holdings of the Federal Reserve, foreign investors, and government accounts, which do not respond to the liquidity premium. Over the 2010s, these exclusions leave thirty-one percent of marketable Treasury debt outstanding, or twenty-two percent of GDP. Finally, the balance-sheet composition B/M is set to clear the bond market at a liquidity premium of $i^b = 2\%$ to target the convenience-yield for near-money assets (Krishnamurthy and Vissing-Jorgensen, 2012).[18] Table 1 summarizes the parameter values and Table 2 compares model and data moments.

[17] While bonds in the model are zero maturity, one can re-interpret them as pure-discount nominal claims that mature into money at a Poisson rate μ and trade at price q below the price of money. Since households rebalance freely between expenditure shocks, allocations depend on the maturity rate and the price discount only through the nominal spread $i^b = \mu(\phi/q - 1)$. Hence, the economy is therefore identical to that of Section 2 along any perfect-foresight transition path and at steady state. The only meaningful difference would be at time zero if an unanticipated shock arrives. The finite-maturity bond price changes on impact and revalues the outstanding stock, while instantly-maturing debt reprices its interest rate instead. I shut this channel down throughout.

[18] The two-percent target is measured at the short end, as the spread between short-term Treasury yields and the own rate on transaction balances. Spreads between money and longer maturities add a term premium that compensates for interest-rate risk, which the model does not price.

Moment	Model	Data	Source
Targeted			
Unemployment rate	0.06	0.06	postwar average
Transaction balances / GDP	0.30	0.25–0.30	M1, post-1980 (Lucas and Nicolini, 2015)
Liquidity premium i^b	2%	$\approx$ 2%	short-rate spread, post-1980 (Ireland, 2009)
Public liabilities / GDP	0.52	0.52	money plus domestically, privately held debt
Interest semi-elasticity	2.1	2.1	post-1980 (Ireland, 2009)

Table 2: Model and data moments.

Comparative statics. Figure 1 illustrates the effect of changes in money growth (inflation) while holding the liquidity premium at its calibrated value. Higher inflation at a fixed premium lowers the real return on money and bonds. Households respond by holding less wealth (top-left panel). Lump-sum transfers rise with inflation via seigniorage revenue (bottom-left), which further incentivizes the reduction in precautionary savings. Money demand falls more slowly than total wealth (top-middle), which reflects that even though the premium is fixed, lower wealth households plan smaller lumpy purchases and so hold less money. Bond demand falls faster than money demand, so sustaining the liquidity premium requires tilting the composition of public assets toward money (top-right panel). Dispersion in consumption rises for both c and y (bottom-middle panel) and the average wedge between the marginal utilities of flow and lumpy consumption widens — the allocation gets further away from first best as inflation rises. Intuitively, despite the fact that the government is providing more insurance through transfers, the inflation tax on households' precautionary wealth buffers is stronger, leading to less consumption insurance.

Figure 2 illustrates the effect of changes in the liquidity premium while holding inflation at two percent. Increasing the liquidity premium with fixed inflation raises r^b and i^b one for one. Households increase their demand for both bonds and money, even though the latter has become relatively more expensive. Supporting a higher premium requires supplying more bonds relative to money (top-right panel). Transfers fall as the real interest bill on the growing bond stock increases. Both dispersion measures fall, and the average marginal-utility gap closes toward one; rising real wealth improves households' self-insurance.

While the two figures only illustrate slices of the set of stationary equilibria along each instrument, they point to where stationary welfare is highest: every real outcome improves as bond returns increase. The stationary, welfare-maximizing policy is $\pi = -2\%$ and $i^b = 1.2\%$, so $r^b \approx 0.94\rho$, but stationary welfare is relatively flat over policies that imply the same bond return.[19] This policy represents a myopic optimum since it ignores transitional dynamics and

[19]Comparable exercises across stationary equilibria appear in Aiyagari and McGrattan (1998) for the supply of government debt and in Rocheteau et al. (2018b) for open-market operations. In the same two-asset environment here, without income risk, Rocheteau et al. (2018a) compute this optimum and find positive inflation and a positive

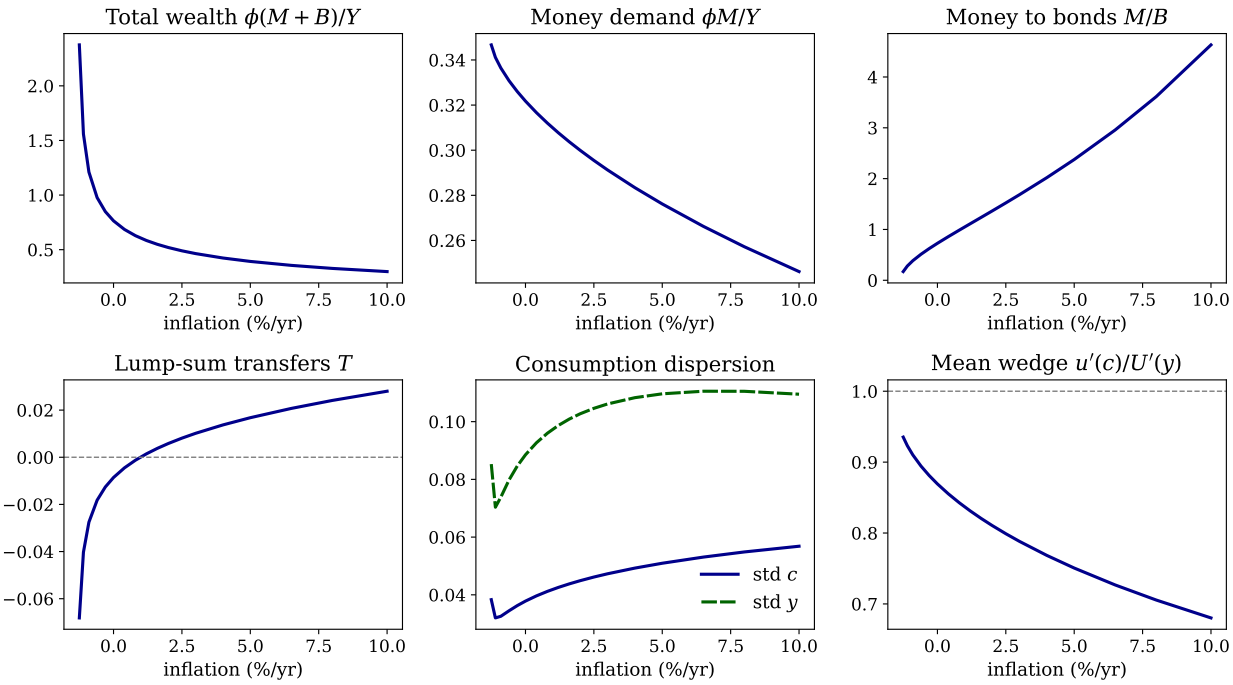


Figure 1: Stationary equilibria across inflation rates, at the calibrated liquidity premium $i^b = 2\%$.

that the optimal stationary policy is not equivalent to the stationary point of an optimal plan. This is the focus of the rest of the paper.

3 Optimal joint monetary-fiscal policy

This section consolidates the Treasury and the central bank into a single planner while Section 4 separates them. The consolidated problem serves as a useful benchmark in which an optimizing central bank has unrestricted remittance power to and from the Treasury, which allows it to independently determine the growth rate of nominal assets and their composition. We will see that separation restricts the set of tools and creates a meaningful trade-off that determines the optimal size of the balance sheet.[20]

I consider cases in which the government can perfectly commit to future policy plans and in which it cannot and have discretion. To do so, I follow the stochastic-replanning approach of Roberds (1987), later developed by Schaumburg and Tambalotti (2007) and applied to heterogeneous-agent economies by Dávila and Schaab (2023). A planner at time t chooses policy $\mathbf{P}_t = \{\dot{M}_s, \dot{B}_s, \mathcal{T}_s(\omega, z)\}_{s \geq t}$ to maximize a utilitarian social welfare function, taking as given the initial measure over nominal wealth and income, $F_t(\tilde{\omega}, z)$, and the initial supply of public assets, $M_t > 0$ and $B_t > 0$. With

premium, with the real bond return just inside the rate of time preference.

[20] With remittances flowing freely between the two authorities, the central bank in this section is really just an accounting partition of the government. One can think of the central bank's balance sheet here as the composition of government liabilities. Its size becomes an independent object only under separation.

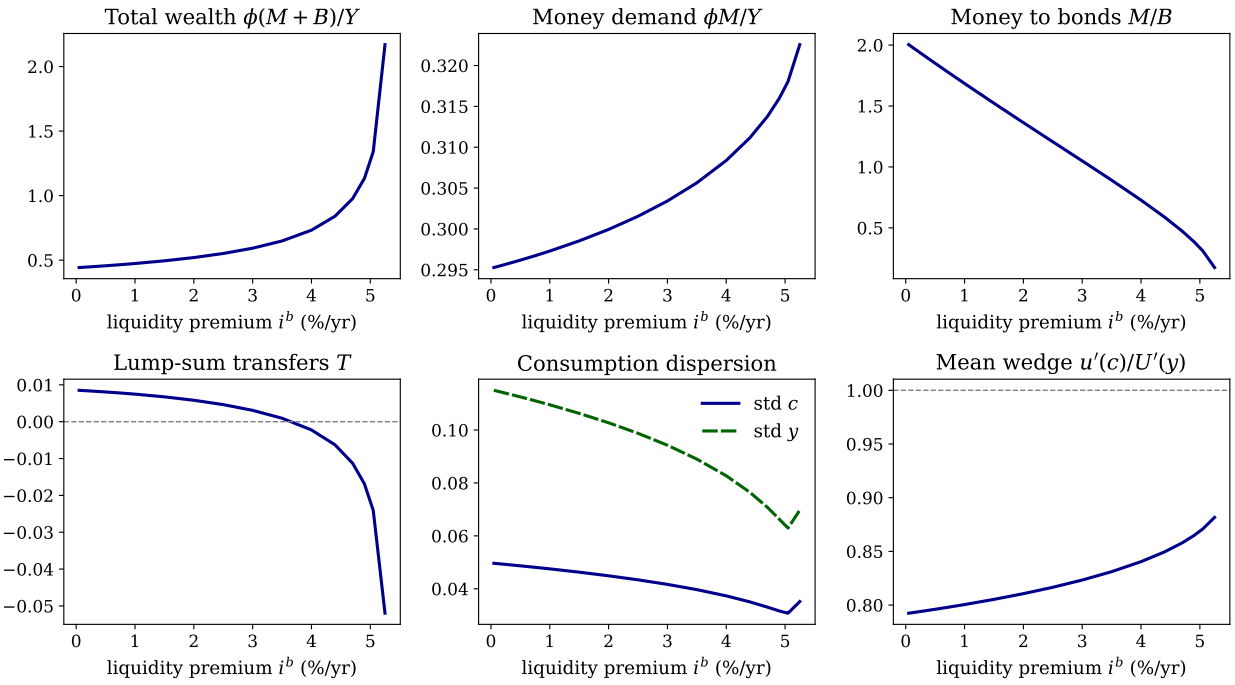


Figure 2: Stationary equilibria across liquidity premia, at the calibrated money growth $g^M = 2\%$.

Poisson intensity $\lambda \geq 0$, a "regime change" occurs: the current planner is replaced by a new planner who re-optimizes, given the state of the economy at that time. The current planner takes into account the probability of being replaced in the future. The case with full commitment is captured by $\lambda \to 0$, while the case with discretion as $\lambda \to \infty$.[21] Finally, I consider two cases that differ in how effectively redistribution can occur through fiscal transfers: i) unconstrained transfers and ii) lump-sum transfers.

3.1 Planner's problem

Let $\mathcal{W}_0(F_0, M_0, B_0)$ denote lifetime discounted social welfare at time 0, and let τ_1 denote the time of the first regime change. The planner inherits the nominal wealth distribution $F_0(\tilde{\omega}, z)$ and the

[21] A subtlety of this formulation is that for intermediate λ, the planner faces genuine aggregate uncertainty about the timing of the regime change, yet the household implementability constraints are derived from a deterministic competitive equilibrium in which households take the policy path as given with certainty. In principle, households should also account for the Poisson arrival of a new planner and the associated potential jump in policy. This inconsistency vanishes at the two limits since under full commitment ($\lambda \to 0$) there is no regime change and hence no aggregate uncertainty, and under full discretion ($\lambda \to \infty$) the Markov-perfect equilibrium has continuous policy paths so households face no jumps. For intermediate λ, a fully consistent formulation would augment the household HJB with a Poisson term $\lambda[W_s^+ - W_s]$, where W_s^+ is the household's value under the successor planner's policy, and analogously augment the planner's backward HJB (26) with $\lambda[\gamma_s^+(x) - \gamma_s(x)]$ and the forward KFE (33) with $\lambda[\psi_s^+(x) - \psi_s(x)]$, where γ_s^+ and ψ_s^+ are the co-states under the successor's optimum.

nominal asset stocks M_0 and B_0. The planner's problem is

$$\mathcal{W}_0(F_0, M_0, B_0) = \max_{\mathbf{X},\mathbf{P}} \mathbb{E}_0 \left[\int_0^{\tau_1} e^{-\rho s} \int \mathcal{U}_s(x)\, dG_s(x)\, ds + e^{-\rho\tau_1} \mathcal{W}_{\tau_1}(F_{\tau_1}, M_{\tau_1}, B_{\tau_1}) \right], \tag{21}$$

subject to (11)–(17), where $x \equiv (\omega, z)$ is the real idiosyncratic state and $\mathcal{U}_s(x) = u(c_s) + \alpha U(y_s)$. The planner's controls are the endogenous equilibrium objects subject to implementability,

$$\mathbf{X} = \{c_t(x), y_t(x), W_t(x), G_t(x), \phi_t, i_t^b\}_{t=0}^{\infty}, \tag{22}$$

together with the policy instruments

$$\mathbf{P} = \{\dot{M}_t, \dot{B}_t, \mathcal{T}_t\}_{t=0}^{\infty}. \tag{23}$$

The nominal asset supplies M_t, B_t are state variables with initial conditions M_0, B_0, evolving through the policy rates $\dot{M}_t, \dot{B}_t$. The planner also chooses the price-of-money path ϕ_t and the liquidity premium i_t^b as equilibrium prices subject to the market-clearing implementability conditions. The real money return $r_t^a = \dot{\phi}_t/\phi_t$ and the real bond return $r_t^b = r_t^a + i_t^b$ are then derived objects, not separate controls. Real money and bond aggregates are $A_t^a \equiv \phi_t M_t$ and $A_t^b \equiv \phi_t B_t$, with total real liabilities $A_t \equiv A_t^a + A_t^b = \phi_t(M_t + B_t)$.

The planner's problem can be re-written in terms of finding a saddle point of the Lagrangian:

$$\mathcal{W}_0 = \max_{\mathbf{X},\mathbf{P}} \min_{\mathbf{M}} \mathbb{E}_0 \left[L(0, \tau_1, F_0, M_0, B_0) + e^{-\rho\tau_1} \mathcal{W}_{\tau_1} \right], \tag{24}$$

where the Lagrangian is

$$
\begin{aligned}
L = \int_{t_1}^{t_2} e^{-\rho s} \Bigg\{ & \int \mathcal{U}_s(x) dG_s(x) \\
& + \Big\{ \psi_s(x) \Big[-\rho W_s + \mathcal{U}_s + \partial_s W_s + \mathcal{A}_s W_s \Big] \\
& + \chi_s(x) \Big[u'(c_s) - \partial_\omega W_s \Big] \\
& + \eta_s(x) \Big[\alpha U'(y_s) - \alpha\, \partial_\omega W_s(\omega - y_s, z) - i_s^b\, \partial_\omega W_s(\omega, z) \Big] \mathbb{1}\{\omega > \underline{\omega}_s(z)\} \\
& + \gamma_s(x) \Big[\mathcal{A}_s^*[G_s] - \partial_s G_s \Big] \Big\} dx \\
& + \theta_s \Big[\dot{M}_s \phi_s + \dot{B}_s \phi_s - i_s^b \phi_s B_s - \int \mathcal{T}_s\, dG_s \Big] \\
& + \mu_s \Big[\int y_s\, dG_s - \phi_s M_s \Big] \\
& + \varphi_s \Big[\int (\omega - y_s)\, dG_s - \phi_s B_s \Big] \Bigg\} ds,
\end{aligned}
\tag{25}
$$

where the real wealth drift is $s_s^\omega = r_s^a \omega + i_s^b(\omega - y_s) + z - c_s + \mathcal{T}_s$ with $r_s^a = \dot{\phi}_s/\phi_s$. The real distribution G_s is the image of the nominal distribution F_s under the price level: with $\omega = \phi_s \tilde{\omega}$, $dG_s(\omega, z) = dF_s(\tilde{\omega}, z)$. The multipliers are $\mathbf{M} = \{\psi_t(x), \chi_t(x), \eta_t(x), \gamma_t(x), \theta_t, \mu_t, \varphi_t\}_{t=0}^{\tau_1}$. In (24)-(25), for any transition time, τ_1, the planner solves a constrained optimal control problem between $[0, \tau_1)$. The Lagrangian in (25) captures the implementability conditions of the planner from the decentralized equilibrium. The second line encodes that the planner's choices induce a particular set of expectations and lifetime welfare for each agent, which determine their individual behavior encoded in lines three and four.[22] The fifth line captures the law of motion of the distribution of wealth and income, while the remaining lines represent the consolidated government budget constraint and the market clearing conditions for money and bonds, respectively. Substituting the household's optimal $a_s = y_s$ into the market clearing conditions, money demand is $\int y_s\, dG_s$ and bond demand is $\int (\omega - y_s)\, dG_s$.

One feature of the saddle-point problem (24)–(25) worth flagging before deriving first-order conditions: the Lagrangian L is integrated over $[0, \tau_1)$ for a *given* realization of the regime-change time τ_1, with the expectation over τ_1 taken outside. Because the implementability constraints

[22] The Euler-y implementability constraint on line four is imposed only in the *interior* (no-depletion) region $\omega > \underline{\omega}_s(z)$, where $\underline{\omega}_s(z)$ is the full-depletion threshold from Lemma 2(i); the integral in this line is restricted to that domain. For $\omega \leq \underline{\omega}_s(z)$, the household's optimality condition for y_s is a Kuhn-Tucker corner solution: the constraint $y_s \leq \omega$ binds, giving $y_s(\omega, z) = \omega$ and a strictly positive multiplier on the constraint. The Euler equation therefore holds as a strict inequality. The indicator $\mathbb{1}\{\omega > \underline{\omega}_s(z)\}$ captures the complementary slackness condition expressed as a domain restriction. The result $y_s = \omega$ for $\omega \leq \underline{\omega}_s(z)$ is imposed into all other terms of the Lagrangian (e.g. the KFE equation, social welfare, etc.).

(the household HJB and FOCs, the KFE, the government budget constraint, and market clearing) must hold at every date s along any sample path, the resulting Euler–Lagrange first-order conditions on $[0, \tau_1)$ are pointwise-in-s conditions that are identical for every realization of τ_1. The Poisson handoff rate λ therefore enters the optimality system only through the terminal and initial conditions at the handoff date (discussed in Section 3.2). This simplicity was first pointed out in Dávila and Schaab (2023) and I leverage it here.

3.2 The optimality system

Taking first-order conditions of the Lagrangian L with respect to the controls $\mathbf{X}$ and the policy instruments $\mathbf{P}$ delivers a coupled forward–backward system that characterizes optimal policy in the two limit cases $\lambda \to 0$ (full commitment) and $\lambda \to \infty$ (full discretion). Proposition 1 states the system in full. The focus of this section is on giving the intuition of each set of equations while the detailed derivations are in the proof in Appendix C.[23]

Proposition 1 (Consolidated optimality system)**.** *For $\lambda \in \{0, \infty\}$, the optimal controls, policy instruments, and multipliers $(\mathbf{X}, \mathbf{P}, \mathbf{M})$ satisfy the equilibrium conditions of Definition 1 together with:*

i) ***Social value.*** *γ_s solves the backward HJB*

$$\rho\gamma_s = \mathcal{U}_s - \theta_s \mathcal{T}_s + \mu_s y_s + \varphi_s (\omega - y_s) + \mathcal{A}_s[\gamma_s] + \partial_t \gamma_s, \tag{26}$$

with terminal condition at the handoff date $\gamma_{\tau_1}(x) = \partial \mathcal{W}_{\tau_1} / \partial F_{\tau_1}(\tilde{\omega}, z)$, where $\tilde{\omega} = \omega / \phi_{\tau_1}$.

ii) ***Revenue shadow price.*** *θ_s satisfies*

$$\theta_s = \partial_\omega \gamma_s + \frac{\psi_s}{g_s} u'(c_s) \quad \text{for all } x, \tag{27}$$

$$\theta_s = \int \partial_\omega \gamma_s \, dG_s + \int \psi_s \, u'(c_s) \, dx \quad \text{under lump-sum transfers,} \tag{28}$$

and evolves according to

$$\dot{\theta}_s = (\rho - r_s^a) \, \theta_s - \mu_s, \tag{29}$$

$$\dot{\theta}_s = (\rho - r_s^b) \, \theta_s - \varphi_s, \tag{30}$$

with terminal condition at the handoff date $\theta_{\tau_1} = -(1/\phi_{\tau_1}) \, \partial \mathcal{W}_{\tau_1} / \partial M_{\tau_1} = -(1/\phi_{\tau_1}) \, \partial \mathcal{W}_{\tau_1} / \partial B_{\tau_1}$.

[23] Appendix C follows the discretize-first method of Dávila and Schaab (2023): time and the idiosyncratic state are discretized and the planner's first-order conditions are taken in the discrete problem, with wealth expressed in detrended-nominal units (relative to the money stock). I then show the limit as the time step and grid mesh $\to 0$ to derive all optimality conditions of the continuous time planner's problem.

iii) ***Household optimality.*** *(χ_s, η_s) satisfy*

$$\chi_s = \frac{\partial_\omega \gamma_s - u'(c_s)}{u''(c_s)}\, g_s, \tag{31}$$

$$\eta_s = \frac{\alpha[\partial_\omega \gamma_s(\omega - y_s, z) - U'(y_s)] + i_s^b[\partial_\omega \gamma_s(\omega, z) - \theta_s]}{\alpha[U''(y_s) + \partial_{\omega\omega} W_s(\omega - y_s, z)]}\, g_s \quad \text{for } \omega > \underline{\omega}_s(z). \tag{32}$$

iv) ***Promise multiplier.*** *Under discretion, $\psi_s \equiv 0$. Under commitment, ψ_s solves the forward KFE*

$$\partial_t \psi_s = \mathcal{A}_s^*[\psi_s] + \partial_\omega \chi_s + \partial_\omega \left[\frac{\alpha\, \eta_s(\omega^+(x), z)}{1 - \partial_\omega y_s(\omega^+(x), z)} \right] + \partial_\omega [i_s^b\, \eta_s], \tag{33}$$

where $\omega^+(x)$ solves $\omega^+(x) - y_s(\omega^+(x), z) = \omega$, with initial condition $\psi_0(x) = 0$.

v) ***Optimal prices.*** *The liquidity premium and the price of money satisfy*

$$\int b_s (\partial_\omega \gamma_s\, g_s + \psi_s\, u'(c_s))\, dx \;=\; \theta_s\, A_s^b \;+\; \int \eta_s\, u'(c_s)\, dx, \tag{34}$$

$$\int \omega\, (\partial_\omega \gamma_s\, g_s + \psi_s\, u'(c_s))\, dx \;=\; \theta_s\, A_s, \tag{35}$$

where $b_s = \omega - y_s$.

Parts i) through iv) of the optimality system in Proposition 1 can be thought of as a set of shadow-price relationships, giving the planner's valuation of each constraint through the multipliers. Part *i*) characterizes the *social value* $\gamma_s(x)$, the welfare value of an additional household at (ω, z). It satisfies a backward HJB in (26) that takes the same form as the household's private value W_s, with the same flow utility $\mathcal{U}_s(x)$, generator $\mathcal{A}_s$, and discount rate ρ. There are three additional terms that capture fiscal externalities: the cost of transferring resources to the household, valued at $-\theta_s\, \mathcal{T}_s(x)$, and the benefits of the household's money and bond demands, valued at $\mu_s\, y_s$ and $\varphi_s\, (\omega - y_s)$, which add to the base for fiscal revenue. The derivative $\partial_\omega \gamma_s$ is the planner's *social marginal value of wealth*—the counterpart of the household's private marginal value of wealth $\partial_\omega W_s = u'(c_s)$.

In part *ii*), the *revenue shadow price* θ_s values a unit of resources in the government budget while the multipliers (μ_s, φ_s) value marginal units of real money and bonds. In (27), when transfers can condition on the state of the household x, the planner equates the shadow value of revenue θ_s to the marginal social value of wealth $\partial_\omega \gamma_s$ plus the value of the utility a marginal unit of resources delivers to household, given by $\psi_s u'(c_s)/g_s$. The term ψ_s is the *promise multiplier*, the shadow value of the lifetime utility implicitly promised to households at (ω, z) by past policy choices and $u'(c_s)$ captures the additional units of lifetime value since the household converts a marginal unit of resources into $\partial_\omega W_s = u'(c_s)$. If transfers are lump-sum, the planner equalizes the shadow value of revenue to the population average marginal social value and promise terms,

in (28).

The co-state equations (29) and (30) price government revenue as an asset, against each liability the government issues. Revenue appreciates at the discount rate net of the liability's real return, $\rho - r_s^a$ against money and $\rho - r_s^b$ against bonds, and pays a "dividend" equal to the social value of the marginal unit outstanding, μ_s for real money and φ_s for real bonds. Subtracting (30) from (29) gives the government's no-arbitrage condition across its own liabilities,

$$\mu_s - \varphi_s = i_s^b \, \theta_s. \tag{36}$$

The social value of a unit of real money exceeds the social value of a unit of real bonds by exactly the fiscal cost of paying interest on bonds.

In part *iii*), the multipliers (χ_s, η_s) price households' private consumption choices through two wedges. The numerator of (31), $\tilde{\chi}_s \equiv \partial_\omega \gamma_s - u'(c_s)$, is the *flow consumption wedge*: the gap between the planner's social marginal value of wealth and the household's private marginal valuation. When $\tilde{\chi}_s > 0$ the planner would like to reduce the household's flow consumption. The opposite holds if $\tilde{\chi}_s < 0$. There is an analogous *lumpy consumption wedge* in the numerator of (32),

$$\tilde{\eta}_s \;\equiv\; \alpha[\partial_\omega \gamma_s(\omega - y_s, z) - U'(y_s)] \;+\; i_s^b[\partial_\omega \gamma_s(\omega, z) - \theta_s] \;+\; \xi_s(\omega, z) \cdot \mathbb{1}_{\{\omega \le \underline{\omega}_s(z)\}}, \tag{37}$$

where $\xi_s(\omega, z) \ge 0$ is the household's Kuhn–Tucker multiplier on the corner constraint $y_s \le \omega$; the indicator extends the wedge to households at $\omega \le \underline{\omega}_s(z)$. It has three parts. The first is the same wedge as in flow consumption, evaluated at the post-depletion wealth $\omega - y_s$ and scaled by the arrival rate α. The second term captures a portfolio wedge that arises because y_s also pins down money holdings ($a_s = y_s$). A marginal increase in y_s raises the household's real money holdings by one unit, valued by the planner at $\partial_\omega \gamma_s - \theta_s$ and scaled by the marginal cost of liquidity i_s^b. The third term, ξ_s, is the corner wedge: zero on the interior region, where the Euler-y condition (13) holds with equality, and strictly positive on the full-depletion region, where $y_s = \omega$ is pinned by the binding constraint.[24] When $\tilde{\eta}_s > 0$, the planner would like to reduce the household's lumpy purchases and, with them, its money holdings.

In part *iv*), the *promise multiplier* $\psi_s(x)$ gives the shadow value of the lifetime utility implicitly promised to households at (ω, z) by past policy choices. Under discretion, $\psi_s(x) = 0$ at all dates

[24] The unified household lumpy KKT condition

$$\alpha \, U'(y_s) \;=\; \alpha \, \partial_\omega W_s(\omega - y_s, z) + i_s^b \, \partial_\omega W_s(\omega, z) + \xi_s, \qquad \xi_s \ge 0, \quad \xi_s \cdot [(\omega) - y_s] = 0$$

nests the interior Euler-y ($\xi_s = 0$) and the corner ($\xi_s > 0$) into a single statement. In the corner region, the implementability multiplier η_s from (32) is not defined (the Euler-y implementability constraint is replaced by the binding policy; see footnote 22), but the corner KKT premium ξ_s still enters $\tilde{\eta}_s$ as the social shadow value of relaxing the constraint.

since the planner is never bound by any past promises. Under commitment, $\psi_s(x)$ satisfies a KFE (33) that evolves with the same generator as the distribution G_s plus three additional terms that give the rate at which promises accumulate at each state, one for each margin the planner distorts: flow consumption, lumpy consumption, and the money–bond portfolio.[25] The planner values changing the promise of utility $W_s(x)$ according to the effect on the multipliers (χ_s, η_s). Since household decision rules depend on the marginal value $\partial_\omega W_s$, an increase in $W_s(x)$ raises $\partial_\omega W_s$ for households just below ω and lowers it for households just above, relaxing the planner's constraint on one side and tightening it on the other. What accumulates then in the value of the promise is the marginal effect on the multipliers. Notice when $\psi_s(x) > 0$, the planner values higher utility promises for households in state x, and when $\psi_s(x) < 0$ they perceive it as a cost.

Part v) of Proposition 1 gives the key conditions that guide optimal policy: the optimality conditions for the liquidity premium and value of money (equivalently, inflation). Equation (34) equates the marginal benefit of a higher liquidity premium to the marginal cost. The marginal benefit, on the left-hand side, is equal to the bond-wealth-weighted marginal social value, $\partial_\omega \gamma_s$, plus marginal promise value, $\psi_s u'(c_s)$. The marginal cost, on the right-hand side, is equal to the fiscal cost $\theta_s A_s^b$ of paying higher interest on outstanding bonds plus the welfare cost $\int \eta_s u'(c_s)\, dx$ of reducing liquidity by tightening the lumpy margin.

The optimal-inflation condition (35) is best read through a devaluation. Consider a marginally lower price of money ϕ_s at date s, holding the nominal wealth distribution F_s and the nominal stocks (M_s, B_s) fixed. Because money is the numeraire for all nominal claims, every household's real wealth $\omega = \phi_s \tilde{\omega}$ scales down in the same proportion, a uniform capital levy on nominal claims. A devaluation dilutes the real value of the government's outstanding liabilities $A_s = \phi_s(M_s + B_s)$, a revenue gain valued at $\theta_s A_s$. The welfare cost is the wealth-weighted marginal social value of the claims destroyed—the left-hand side of (35), denoted $R_s \equiv \int \omega \left(\partial_\omega \gamma_s \, g_s + \psi_s \, \partial_\omega W_s\right) dx$ in what follows. At the optimum the two offset.

Boundary conditions. In time, the terminal conditions equate each backward co-state to the successor's valuation of the state it prices. The two expressions for θ_{τ_1} coincide because the consolidated successor freely re-mixes its liability composition. The initial condition $\psi_0 = 0$ holds because date-0 household values are not inherited. This is the only place the replanning rate λ enters. The interior conditions are the same under commitment and discretion.[26] In terms of

[25] This interpretation connects to the recursive contracts literature (Marcet and Marimon, 2019): the planner carries utility promises as state variables, the envelope condition on those promises is the promise-keeping constraint, and the co-state on that constraint is the shadow price of the promise. Here $W_s(x)$ plays the role of the promise—the continuation value implicitly delivered to household (ω, z) by past policy choices—and the HJB equation is its law of motion. The multiplier $\psi_s(x)$ is therefore the shadow price of that promise.

[26] In the long-run limit $\tau_1 \to \infty$ used for the stationary analysis, the terminal matching is replaced by a transversality requirement that the discounted co-states not explode—in particular $e^{-\rho s}\theta_s A_s \to 0$ at the optimum—selecting the bounded, saddle-path solution.

boundary conditions over wealth, the conditions in Proposition 1 are stated for the interior of the state space. At the borrowing constraint, the same logic applies with the mass point density $m_s(z)$ in place of g_s, aggregate conditions integrate over both, and the adjoint equations carry boundary conditions at $\omega = 0$. These are derived carefully in Appendix C.

3.3 Optimal consolidated policy with unrestricted transfers

First, consider optimal policy with unrestricted transfers $\mathcal{T}_t(\omega, z)$, in which the planner can condition the flow transfer on each agent's idiosyncratic state (ω, z). The following Lemma shows that the multiplier on utility promises vanishes for either $\lambda \to 0$ or ∞, and the planner can achieve the first-best allocation.

Lemma 4 (Commitment premium vanishes under unrestricted transfers)**.** *Under full discretion ($\lambda \to \infty$) or full commitment ($\lambda \to 0$), if transfers are unrestricted then $\psi_t(x) = 0$ for all (x, t), and the transfer first-order condition* (27) *reduces to*

$$\theta_t = \partial_\omega \gamma_t(x) \quad \textit{for all } (x, t). \tag{38}$$

When transfers are unrestricted, the planner equates the marginal social gain of transfers across agents. The proof shows that this implies the constraints on utility promises are completely relaxed for all time and for all agents, $\psi_t(x) = 0$. This holds not only under full discretion $\lambda \to \infty$, where each instantaneous planner has no past promise obligations by assumption, but also under full commitment $\lambda \to 0$. Intuitively, if the planner can equate marginal social values across agents through transfers, terms that differentiate promises across agents in (33) collapse to zero.

By equating the marginal private value of wealth $\partial_\omega W_t$ across the state space, the fiscal side of policy removes every redistributive motive for policy—in both flow consumption c_t and lumpy consumption y_t. The monetary side then corrects the one margin transfers cannot reach: the common distortion in the level of lumpy expenditure. Since the lumpy purchase must be financed from pre-accumulated real balances, $y_t \leq \omega_t$, a liquidity cost ($r_t^a < \rho$ and $i_t^b > 0$) depresses lumpy expenditure uniformly across agents; eliminating it requires both the Friedman rule ($r_t^a = \rho$) and a zero liquidity premium ($i_t^b = 0$).

Proposition 2 (Policy under unrestricted transfers)**.** *Under unrestricted transfers and either full discretion ($\lambda \to \infty$) or full commitment ($\lambda \to 0$), the optimal policy sets $B_t = 0$ and $M_t > 0$ with the real return on money at the Friedman rule, $r_t^a = \rho$. The liquidity premium is zero ($i_t^b = 0$) and the allocation is time and state invariant, coinciding with the first-best in* (2)–(3)*.*

Two aspects of this result deserve emphasis. First, bonds are redundant at the Friedman rule. With $r_t^a = r_t^b = \rho$ ($i_t^b = 0$), money and bonds offer the same real return, and households are no longer penalized for holding wealth in money. Any small positive bond supply in the hands

the public would also clear markets at $i_t^b = 0$, so setting $B_t = 0$ is therefore without loss but not unique. Second, perfect insurance over *flow* consumption does not eliminate the demand for money. Even when $u'(c_t) = \theta_t$ is equalized across agents, households still need to accumulate liquid wealth to finance lumpy purchases. The fiscal side cannot undo this friction since transfers are not allowed to be lumpy. Policy distorts this margin twice. A positive liquidity premium distorts the composition of wealth, steering savings toward bonds. A real return on money below the rate of time preference distorts the level of wealth, even at $i_t^b = 0$. Removing both distortions requires $r_t^a = \rho$, and since the bond return cannot exceed the rate of time preference in equilibrium, the premium goes to zero.[27]

In terms of a central bank balance sheet, its size is defined only through what assets it leaves in public hands. A larger balance sheet holds more of the debt and issues more money against it, so bonds held by the public fall and the real supply of money rises. The benchmark pushes this exchange to its limit. The optimum drives the public's bond holdings low enough to eliminate the liquidity premium and supplies high enough real money balances, financed by deflation at the rate of time preference.

3.4 Optimal consolidated policy with discretion and lump-sum transfers

When transfers are restricted to be lump sum, any redistributive effect depends on how the transfers are financed. Inflation is a proportional tax on all nominal claims, redistributing from wealth-rich to wealth-poor households. The liquidity premium taxes wealth according to its composition, falling on the share held in money. In this section, I study the stationary optimum under full discretion, where the two taxes map one-to-one into the planner's two instruments. Both nominal stocks must grow at the common rate $-r_{ss}^a$ for real wealth to remain stationary, so the growth rate sets the inflation tax; the balance sheet composition M_{ss}/B_{ss} sets the liquidity premium through bond-market clearing.

Under full discretion ($\lambda \to \infty$), the multiplier on utility promises collapses to $\psi_t(x) \equiv 0$ at every instant—each planner is only constrained by the state it inherits at time t, (F_t, M_t, B_t), but not by past promises. Under lump-sum transfers, the marginal cost of government revenue θ_t equals the average social marginal value of wealth, (28) becomes

$$\theta_t = \int \partial_\omega \gamma_t(x)\, dG_t(x). \tag{39}$$

A lump-sum transfer reaches every household equally, so revenue is valued at the population average. Setting $\psi_{ss} \equiv 0$ in the two pricing conditions of Proposition 1, the optimal-inflation

[27] Unrestricted transfers also restore the *feasibility* of the Friedman rule. With lump-sum transfers, aggregate wealth diverges as $r_t^a \to 0^+$ due to the precautionary savings motive. With unrestricted transfers, the planner controls the savings drift $s_t^\omega(x)$ pointwise and can pin aggregate wealth at a finite level even at $r_t^a = \rho$.

condition reads $\int \omega\, \partial_\omega \gamma_{ss}\, dG_{ss} = \theta_{ss} A_{ss}$ and the bond-rate condition $\int (\omega - y_{ss})\, \partial_\omega \gamma_{ss}\, dG_{ss} = \theta_{ss} A^b_{ss} + \int \eta_{ss}\, u'(c_{ss})\, dx$. Each compares a wealth-weighted marginal social value of wealth against θ_{ss} times the corresponding asset stock. With θ_{ss} the unweighted average, each condition is a statement about a covariance.

Assumption 2 (Bounded relative risk aversion)**.** *There exists $\bar{\gamma} < \infty$ such that $-cu''(c)/u'(c) \leq \bar{\gamma}$ for all $c > 0$. In addition, absolute risk aversion $-u''(c)/u'(c)$ remains finite as $c \downarrow 0$.*

Proposition 3 (Optimal consolidated policy under full discretion)**.** *Under Assumptions 1 and 2, full discretion, and lump-sum transfers, the optimal steady state has $r^{a*}_{ss} < \rho$, and the optimal-inflation and bond-rate conditions (35) and (34), with $\psi_{ss} \equiv 0$, take the covariance form*

$$\mathrm{Cov}_{G_{ss}}(\omega,\, \partial_\omega \gamma_{ss}) = 0, \tag{40}$$

$$\mathrm{Cov}_{G_{ss}}(\omega - y_{ss},\, \partial_\omega \gamma_{ss}) = \int \eta_{ss}\, u'(c_{ss})\, dx, \tag{41}$$

where $\mathrm{Cov}_{G_{ss}}$ is the cross-sectional covariance under G_{ss} and η_{ss} is the lumpy implementability multiplier (32).

The two conditions share a similar structure: optimal consolidated policy targets the covariance between an asset position and the social marginal value of wealth $\partial_\omega \gamma_{ss}$. The inflation tax falls on total wealth ω, so (40) sets the optimal r^{a*}_{ss} where $\mathrm{Cov}_{G_{ss}}(\omega, \partial_\omega \gamma_{ss}) = 0$: when wealthier households have a lower social value of wealth the covariance is negative and the planner wants more inflation, and when higher it wants less. The optimum balances at zero. The liquidity premium instead falls on bond holdings $\omega - y_{ss}$, so (41) is the same trade-off for i^{b*}_{ss}—except that raising the premium also tightens the household's lumpy implementability constraint, so the optimum equates the bond-holdings covariance to that implementability cost $\int \eta_{ss} u'(c_{ss})\, dx$ rather than to zero. Assumption 2 is a standard preference regularity condition (Achdou et al., 2022): it keeps the consumption policy increasing in wealth.

The sign of each covariance depends on how the social marginal value of wealth varies across the distribution. Substituting the stationary co-state values $\mu_{ss} = (\rho - r^a_{ss})\theta_{ss}$ and $\varphi_{ss} = (\rho - r^b_{ss})\theta_{ss}$ from (29)–(30), together with the stationary transfer $T_{ss} = -r^a_{ss} A_{ss} - i^b_{ss} A^b_{ss}$ implied by the government budget constraint, where $A_{ss} \equiv \int \omega\, dG_{ss}$ is total real wealth, we can write the social value of wealth (26) as

$$\rho \gamma_{ss} = \mathcal{U}_{ss} + \rho \theta_{ss} A_{ss} + \underbrace{(\rho - r^a_{ss})\theta_{ss}[\omega - A_{ss}]}_{\text{inflation-tax redistribution}} - \underbrace{i^b_{ss}\, \theta_{ss}[(\omega - y_{ss}) - A^b_{ss}]}_{\text{liquidity-premium redistribution}} + \mathcal{A}_{ss}[\gamma_{ss}], \tag{42}$$

where $A^a_{ss} = \phi_{ss} M_{ss} = \int y_{ss}\, dG_{ss}$ is total real money holdings and $A^b_{ss} = \phi_{ss} B_{ss} = \int (\omega - y_{ss})\, dG_{ss}$ is total real bond holdings. The two underlined terms capture the net fiscal externality of the

household, relative to the mean. The inflation-tax redistribution term $(\rho - r^a_{ss})\theta_{ss}[\omega - A_{ss}]$ is positive when $\omega > A_{ss}$: a household with above-average total wealth is a net contributor to the government's inflation-tax revenue. The inflation tax erodes their real balances at rate $(\rho - r^a_{ss})$, and this revenue is redistributed uniformly via lump-sum transfers. Since the household pays more in inflation tax than they receive in transfers, adding them improves the government's fiscal position, which the planner values at shadow price θ_{ss}. Conversely, a household with below-average wealth is a net recipient of transfers—adding them tightens the government's budget. The liquidity-premium term $-i^b_{ss}\theta_{ss}[(\omega - y_{ss}) - A^b_{ss}]$ captures an analogous fiscal externality from bond interest payments: a household with above-average bond holdings $\omega - y_{ss}$ receives more interest from the government than the average, imposing a fiscal cost. The term $\rho\theta_{ss}A_{ss}$ is a constant that captures the average social benefit of aggregate real wealth.

What forces stand behind the covariances in (40) and (41)—and hence behind the temptation to inflate and the temptation to tax liquidity—is clearest after decomposing the social marginal value of wealth.

Lemma 5 (Decomposition of the social marginal value of wealth)**.** *At any stationary equilibrium of the consolidated problem, the social marginal value of wealth satisfies*

$$\partial_\omega \gamma_{ss} \;=\; u'(c_{ss}) \;+\; \theta_{ss}\, \partial_\omega \gamma^f_{ss}, \tag{43}$$

where $\gamma^f_{ss}(x) = \mathbb{E}\int_0^\infty e^{-\rho s}[(\rho - r^a_{ss})\, y_s + (\rho - r^b_{ss})\,(\omega_s - y_s)]\, ds$ *is the expected discounted stream of revenue the government will collect from the household's future asset positions.*

The social marginal value of wealth can be decomposed into the household's private marginal value of wealth, equal to $\partial_\omega W = u'(c)$, plus the marginal fiscal value of its wealth, $\theta\partial_\omega\gamma^f$. A marginal unit of wealth flows into future money and bond holdings, on which the household forgoes the returns $\rho - r^a_{ss}$ and $\rho - r^b_{ss}$ relative to the planner's discount rate—the user costs of money and of bonds. Since both assets are government liabilities, the forgone returns are government revenue valued at the shadow price θ_{ss}.[28]

Substituting (43) into (40), the optimal-inflation condition becomes a balance of two channels,

$$\mathrm{Cov}_{G_{ss}}(\omega,\, u'(c_{ss})) \;+\; \theta_{ss}\,\mathrm{Cov}_{G_{ss}}(\omega,\, \partial_\omega \gamma^f_{ss}) \;=\; 0. \tag{44}$$

The first term is the redistribution channel. Consumption rises with wealth, so marginal utility falls with it and the covariance is negative. Taxing nominal wealth and rebating the proceeds uniformly moves resources toward households whose marginal utility is high. This is the temptation to inflate. The second term is the fiscal-value channel, and it is what restrains

[28]The proof (Appendix A.8) uses only the linearity of the social-value equation (42), so the decomposition holds under discretion and commitment alike, with the regime determining only the value of θ_{ss}.

the temptation. Since wealth compounds, a marginal unit given to a higher wealth household keeps generating user-cost revenue. The fiscal value $\partial_\omega \gamma^f_{ss}$ therefore tends to rise with wealth and the second covariance in (44) tends to be positive. A stationary optimum exists exactly where the fiscal-value channel grows strong enough to absorb the redistributive pull. Whether any stationary policy achieves this balance is a quantitative question, which I turn to next.

Quantifying the temptation to inflate. Figure 3 illustrates the two channels of (44) at the calibrated steady state of Section 2.5. The left panel plots the marginal social value of wealth and its decomposition across the wealth distribution. Private marginal utility falls steeply with wealth while the fiscal value rises with wealth. The fiscal value is an order of magnitude lower than private marginal utility across the wealth distribution, so private marginal utility comprises most of the social marginal value of wealth. The right panel plots each household's contribution to the marginal benefit and cost of inflation. The red-dashed line shows the redistributive gain —positive below average wealth and negative above—and it dominates the fiscal cost of higher inflation. The gain integrates to +0.044, while the fiscal cost integrates to +0.010, so the net marginal gain from inflation is +0.034 at the calibrated policy. At two percent inflation and the observed liquidity premium, the discretionary planner gains from a marginal increase in the inflation tax.

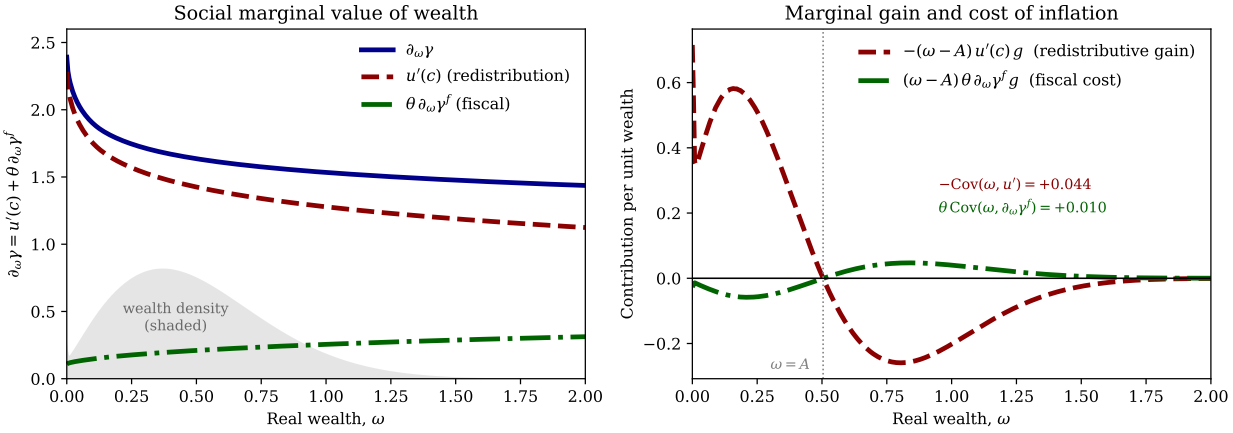


Figure 3: The temptation to inflate at the calibrated steady state.

Quantifying the temptation to increase the liquidity premium. Figure 4 repeats the exercise for the optimal steady-state bond-rate condition (41). The left panel plots the household's portfolio choice by wealth. Money holdings flatten quickly with wealth, so bond holdings absorb nearly all wealth growth above a certain level. The premium is therefore a payment to households who look much like the inflation tax's contributors.

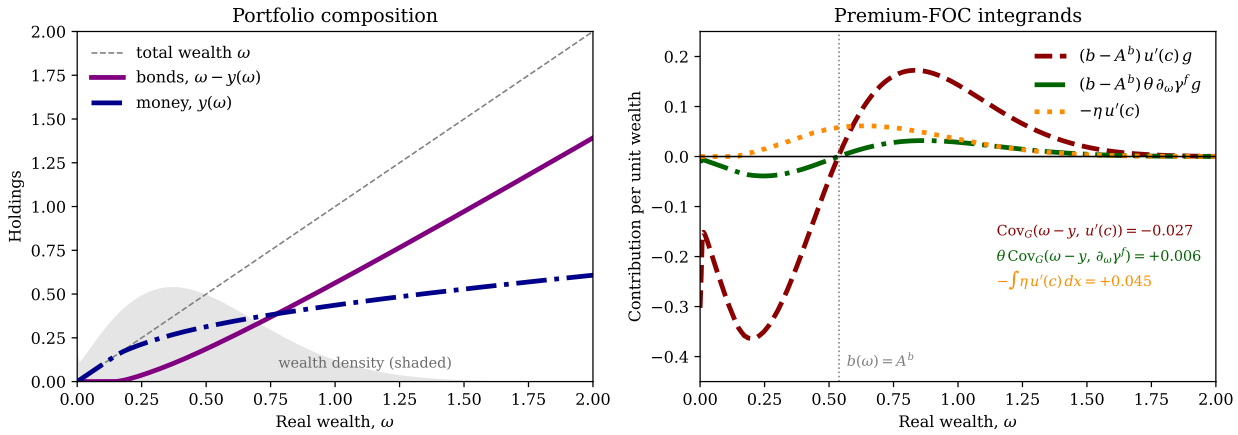


Figure 4: The temptation to tax liquidity at the calibrated steady state.

The right panel of Figure 4 plots the three integrands of (41). The dashed-red line illustrates the redistributive cost of raising i^b, by wealth. Increasing the liquidity premium taxes high marginal-utility, low-wealth households and subsidizes bond holdings of low marginal-utility, high-wealth households. This term integrates across the wealth distribution to -0.027. The dash-dotted green line shows the fiscal value of increasing the liquidity premium. Households with below-average bond holdings receive less in premium payments than the per-capita fiscal cost they bear, so a higher premium drains their wealth and, with it, the future revenue the government would have collected on that wealth. For higher wealth, this flips and increasing the liquidity premium improves the fiscal value of the high-wealth households. This term integrates to +0.006. The dotted-orange line shows the liquidity effect of increasing the premium. Retained wealth keeps generating user-cost revenue, and households do not value this revenue when they spend, so the planner prefers less spending than households choose ($\tilde{\eta} > 0$). A higher premium raises the cost of carrying the money that funds this spending, moving every household's choice in the planner's preferred direction; the multiplier η_{ss} translates the wedge through the household's own first-order condition, whose concavity gives it the opposite sign, so $-\eta_{ss} u'(c_{ss})$ is positive household by household. The liquidity term integrates to +0.045. Collectively, the three areas sum to +0.024. In sum, keeping inflation fixed, the discretionary planner wants a higher liquidity premium.

Optimal time-consistent hyperinflation. Figure 5 evaluates both optimality conditions as steady-state money growth (left-side panels) or the liquidity premium (right-side panels) change, keeping the other fixed.[29] The top-left panel plots the redistributive gain and the fiscal cost of inflation. The gain exceeds the cost even as steady-state inflation reaches 25%. While the redistributive ben-

[29] Lump-sum transfers always change to balance the government budget constraint. For the right-side panels, the mix B/M is adjusted to target a particular liquidity premium.

efit falls with higher inflation, so does the fiscal cost since the wealth distribution starts to shift left and narrow. The two lines never cross: extending even to 150% inflation, the gain exceeds the cost at every point. The bottom-left panel shows what this pursuit costs. Stationary welfare falls monotonically with inflation even as the planner's condition points toward more of it, because the condition prices the initial revaluation of nominal wealth that comparisons across stationary equilibria omit.

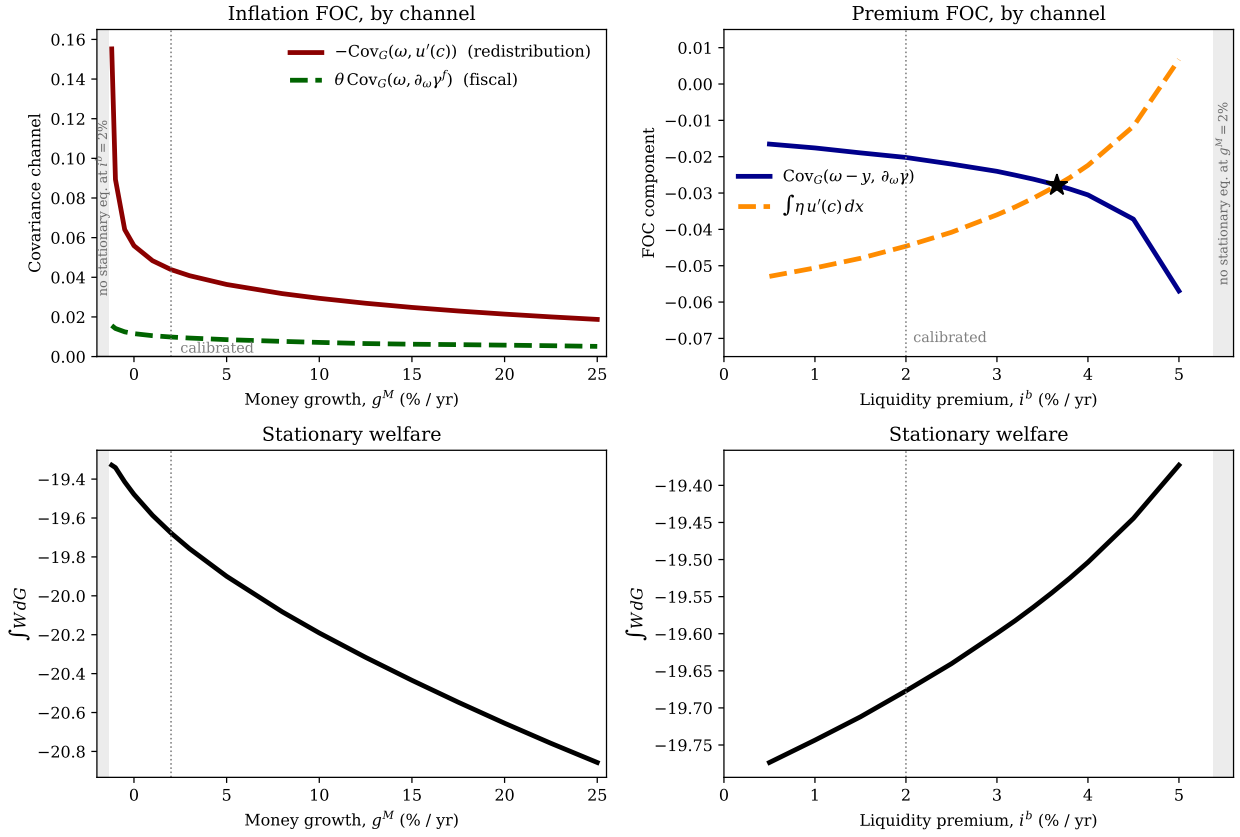


Figure 5: The optimal-inflation and liquidity-premium conditions across stationary equilibria.

The top-right panel repeats the exercise for the optimal liquidity-premium condition. Keeping inflation fixed, the optimal liquidity premium is $i^{b*}_{ss} = 3.7\%$, above the calibrated steady-state value of two percent. Beyond that point, the redistributive costs outweigh the fiscal and liquidity benefits. If the planner were prevented from running inflation high, the optimal time-consistent mix of money and bonds should be $B_{ss}/M_{ss} = 1.2$, tilted further toward bonds than the mix in the data. The inflation margin is where discretion fails.

3.5 Optimal consolidated policy with commitment and lump-sum transfers

The central difference between commitment and discretion is the promise multiplier ψ_t. Households build their wealth in anticipation of the returns policy will deliver. A committed planner

internalizes this, and the returns it commits to are what sustain precautionary saving. The promise multiplier records the value of honoring this obligation. Figure 6 illustrates the promise multiplier at the calibrated steady state, across the distribution of wealth. As in Section 3.4, I focus on the stationary optimum.

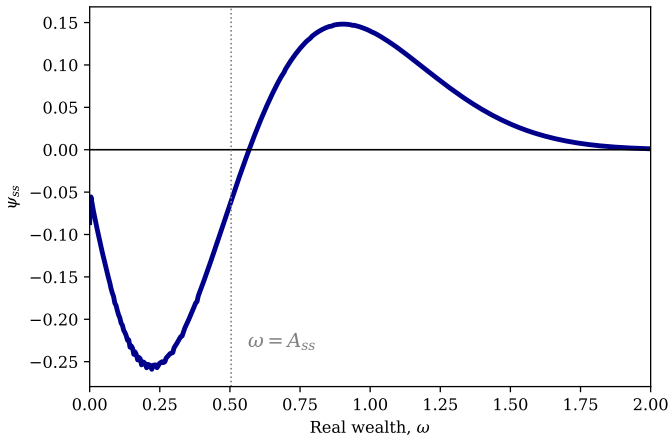


Figure 6: The promise multiplier at the calibrated steady state.

The promise multiplier is negative below average wealth and positive above. Promised utility at high wealth is the reward that sustains savings, so the planner values honoring it, $\psi > 0$, since the savings it sustains generate revenue. Conversely, promised utility at low wealth acts as insurance against not accumulating, so at the margin it dulls precautionary saving, reduces the tax base, and its shadow value is negative, $\psi < 0$.

The promise multiplier affects both the marginal benefits and costs of inflation and the liquidity premium in (34) and (35). Reducing inflation subsidizes nominal wealth and increases promised utility, worth $\int \omega\, \psi_{ss} u'(c_{ss})\, dx$ to the planner. However, the lost revenue reduces transfers (or increases taxes), worth $\theta_{ss} A_{ss}$. From (28), $\theta_{ss} = \int \partial_\omega \gamma_{ss}\, dG_{ss} + \int \psi_{ss}\, u'(c_{ss})\, dx$, where the second term is the promise effect of the change in transfers. Proposition 4 nets them into a single premium on each margin. The bond-rate condition carries the same structure with bond holdings $\omega - y_{ss}$ in place of ω.

Proposition 4 (Optimal consolidated policy under full commitment)**.** *Under Assumptions 1 and 2, full commitment, and lump-sum transfers, the stationary equilibrium satisfies* $r_{ss}^{a,\text{comm}} < \rho$*, and the optimal-inflation and bond-rate conditions are:*

$$\text{Cov}_{G_{ss}}(\omega,\, \partial_\omega \gamma_{ss}) + \Psi_{ss} = 0, \tag{45}$$

$$\text{Cov}_{G_{ss}}(\omega - y_{ss},\, \partial_\omega \gamma_{ss}) + \Psi^b_{ss} = \int \eta_{ss}\, u'(c_{ss})\, dx, \tag{46}$$

where the commitment premiums are given by $\Psi_{ss} = \int (\omega - A_{ss})\, \psi_{ss}\, u'(c_{ss})\, dx$ *for total wealth and* $\Psi^b_{ss} =$

$\int((\omega - y_{ss}) - A^b_{ss})\,\psi_{ss}\,u'(c_{ss})\,dx$ *for bond wealth, and* ψ_{ss} *solves the forward KFE* (33) *with* $\partial_t \psi_s = 0$.

Proposition 4 identifies the optimal stationary inflation (via $\dot{M}_{ss} = \dot{B}_{ss} = -r^a_{ss}$) and balance-sheet ($M_{ss}/B_{ss}$) policies under commitment. Each condition is identical to its discretion counterpart in Proposition 3, but adds a commitment premium—Ψ_{ss} for inflation, Ψ^b_{ss} for the liquidity premium—that the discretion planner sets to zero. Each premium is itself a covariance. For example, $\Psi_{ss} = \mathrm{Cov}_{G_{ss}}(\omega, \psi_{ss}u'(c_{ss})/g_{ss})$, and Ψ^b_{ss} is the same with bond holdings $\omega - y_{ss}$ replacing ω.

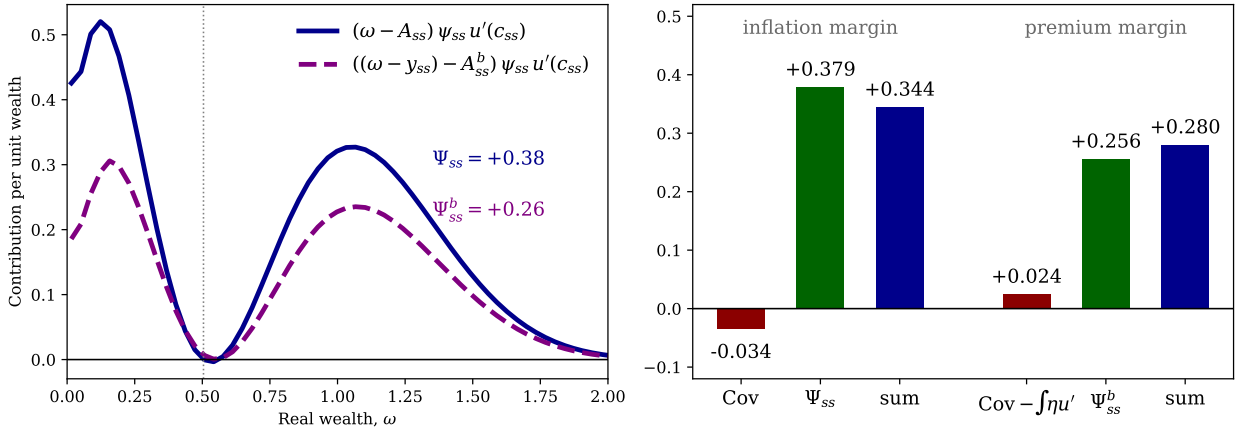


Figure 7: The commitment premium at the calibrated steady state.

Figure 7 illustrates how the promise multipliers convert into the commitment premiums. The left panel plots the integrands of Ψ_{ss} and Ψ^b_{ss}. Both integrands are positive at every wealth level, with twin peaks on either side of average wealth. A marginal increase in inflation shifts delivered utility from households with above average wealth to those with below average wealth—it breaks valuable promises and enriches costly ones—so both tails contribute positively to the premium, and the contribution vanishes at average wealth. The right panel computes the premiums against the redistributive terms of the two optimality conditions. With respect to inflation, the premium Ψ_{ss} is roughly eleven times the redistributive covariance, so the condition fails in the opposite direction from discretion. Honoring promises makes marginal inflation costly, and the committed planner is pulled toward deflation and higher returns to wealth. With respect to the liquidity premium, the promise channel adds to the discretionary push toward a higher premium.

The commitment premium Ψ_{ss} can be characterized further by re-writing the individual promise multipliers $\psi_{ss}(x)$ that comprise it. These solve the forward equation (33), and the implementability multipliers, $\chi_{ss}(x)$ and $\eta_{ss}(x)$, drive it. As a household saves, spends, and switches income states, the multipliers it carries accumulate, and $\psi_{ss}(x)$ records the total amount accumulated along the histories that bring households to state x. The promise multipliers are therefore built up from the implementability multipliers, and those are scaled versions of the wedges. It

is helpful to first re-write the wedges, using the decomposition of Lemma 5, as functions of the fiscal value of a household's wealth, $\partial_\omega \gamma^f_{ss}$, alone.

Lemma 6 (The wedges price fiscal externalities)**.** *At any stationary equilibrium of the consolidated problem, the flow consumption wedge and, on the interior region $\omega > \underline{\omega}_{ss}(z)$, the lumpy consumption wedge satisfy*

$$\tilde{\chi}_{ss} \;=\; \theta_{ss}\,\partial_\omega \gamma^f_{ss}, \qquad \tilde{\eta}_{ss} \;=\; \theta_{ss}\Big[\alpha\,\partial_\omega \gamma^f_{ss}(\omega - y_{ss}, z) \;-\; i^b_{ss}(1 - \partial_\omega \gamma^f_{ss}(\omega, z))\Big]. \tag{47}$$

Every private marginal-utility term cancels from the wedges. The first identity is immediate from Lemma 5; the second uses the household's lumpy Euler (13) (Appendix A.10). What remains is the fiscal externality alone. The planner disagrees with a household's consumption-savings choices only because saved wealth generates future fiscal revenue—worth θ_{ss} per unit—that the household does not internalize. In the flow consumption wedge, the planner values marginal wealth above the household by exactly the fiscal value $\partial_\omega \gamma^f_{ss}$. In the lumpy wedge, two fiscal effects compete. Spending y_{ss} at an expenditure shock destroys post-shock wealth and the revenue attached to it (the first term, arriving at rate α). Holding the marginal unit in money rather than bonds saves the government interest costs (the second term).

The promise multiplier in (33) can then be read as the accumulation, along household histories, of the wealth-gradients $\partial_\omega \chi_{ss}$ and $\partial_\omega \eta_{ss}$—how the fiscal externalities carried by a household's spending margins change with its wealth. The next result says that the planner does not need to track this accumulation over time and across households. It is enough to know two objects for each household: the fiscal externality the household carries today, and how one more unit of wealth would move its standing in the redistribution the planner engineers — what I refer to as the household's entry in a redistributive ledger.

First, collect the implementability multipliers into a single term, $\Phi_{ss} \equiv \chi_{ss} + [\alpha\eta_{ss}/(1 - \partial_\omega y_{ss})](\omega^+(x), z) + i^b_{ss}\,\eta_{ss}$, with $\omega^+(x)$ the pre-depletion wealth. Second, define $q(x) \equiv (\omega - A_{ss})\,u'(c_{ss}) - \mathrm{Cov}_{G_{ss}}(\omega, u'(c_{ss}))$, which records a household's marginal-utility-weighted exposure to the inflation tax, net of the uniform rebate, and expressed relative to the population average — their entry in a 'redistributive ledger'. Wealth-rich households carry positive entries in the ledger. They lose from marginal inflation and are its net contributors. Wealth-poor households carry negative entries and are its net recipients. Like any ledger, it balances—the entries integrate to zero across the population, $\int_x q(x) dG_{ss}(x) = 0$. Let $\tilde{q}_{ss}(x) \equiv \mathbb{E}_x \int_0^\infty q(x_s)\,ds$ denote a household's entry cumulated over its expected future path. It solves the Poisson equation $\mathcal{A}_{ss}[\tilde{q}_{ss}] = -q$ on the household's stationary generator.[30] The wealth-gradient $\partial_\omega \tilde{q}_{ss}$ then answers a simple ques-

[30]The integral converges because entries integrate to zero under the stationary distribution, and $\tilde{q}_{ss}$ solves the Poisson equation uniquely up to an additive constant that drops from everything that follows; see the proof of Lemma 7 in Appendix A.11.

tion. Suppose the household received one more unit of wealth today. Over its remaining lifetime, would its losses to the inflation tax, net of the rebates it receives and valued at its marginal utility, be larger?

Lemma 7 (A dual representation for the commitment premium)**.** *The commitment premium of Proposition 4 satisfies*

$$\Psi_{ss} \;=\; -\,\theta_{ss} \int \partial_\omega \tilde{q}_{ss}(x)\, \Phi_{ss}(x)\, dx. \tag{48}$$

In (48), the premium pairs the planner's pressure for more saving, Φ_{ss}, with the additional lifetime tax contributions that an extra unit of wealth brings, $\partial_\omega \tilde{q}_{ss}$. The premium leans against the redistributive pull toward inflation. The reason is that the tax base is endogenous. Households accumulate against the return they anticipate, so anticipated inflation means under-accumulation of exactly the wealth the planner is counting on. The promise multiplier records this margin, and it is concentrated on exactly the households who built wealth against the promised return. Inflating collects the base but breaks the promises that make the base accumulate. Only the committed planner prices the second effect, and it tilts policy toward a higher anticipated return on nominal assets.

Existence of a stationary optimum under commitment As alluded to in Section 3.4, under discretion the redistributive gain from inflation exceeds its fiscal cost at every stationary monetized equilibrium. Left to itself, consolidated discretion inflates the economy toward demonetization. Commitment reverses the direction into a push toward the Friedman rule. The push runs up against a tax limit. Wealth demand grows as the gap $\rho - r^b$ closes, and the lump-sum tax that finances the deflation rises with it, toward the income of the poorest households. The planner weighs that tax by their marginal utility, so the pull toward the Friedman rule cannot survive all the way to the limit. A stationary optimum is then the point at which the deflationary pull and the cost of the tax balance out, equal to $\pi^* = -3.35\%$ and $i^{b*} = 0$ in the calibrated economy illustrated in Figure 8.[31]

Taken together, the two consolidated regimes resolve the tension between financing social insurance and incentivizing self-insurance at opposite corners. Discretion taxes the buffer stock all the way to demonetization. Commitment protects it nearly all the way to the Friedman rule, surrendering the financing instead.

[31] Wealth demand diverges slowly as the bond return limits toward the rate of time preference. Households spend wealth on lumpy expenditure at a rate that rises with wealth, which tempers the standard incomplete-markets divergence. At the baseline calibration the balancing point sits two hundredths of a percentage point from the Friedman rule, where stationary wealth reaches eleven times output and the tax reaches three quarters of the income of the poorest households.

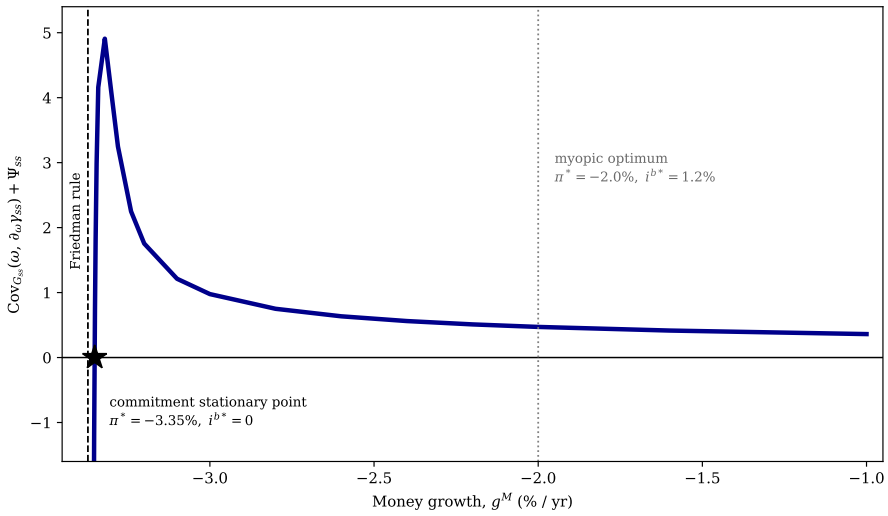


Figure 8: The commitment optimal-inflation condition along the deflation path, at the premium floor.

4 Optimal central bank balance sheets

I now turn to the case where monetary and fiscal policy are separate, each with its own policy instruments and budget constraints. Specifically, I assume the central bank controls its monetary liabilities, assets, and remittances, $(\dot{M}_t, \dot{B}_t^{CB}, R_t)$, subject to (6), while the Treasury sets transfers and debt issuance $(T_t, \dot{B}_t^T)$, subject to (5).

To study the central bank's problem in isolation I impose two restrictions: i) the Treasury follows a passive fiscal rule and ii) the central bank faces constraints on remittances. First, following Leeper (1991), the Treasury sets lump-sum transfers by a debt-stabilizing rule:

$$T_t = \bar{T} + \kappa \left(\bar{b} - \phi_t B_t^T \right), \tag{49}$$

where $\bar{T}$ is the intercept of the rule, $\bar{b}$ is a target for the real value of outstanding government debt, and $\kappa \geq 0$ governs the speed of fiscal adjustment. When real debt $\phi_t B_t^T$ rises above its target $\bar{b}$, the Treasury lowers transfers, raising the primary surplus, and reducing real debt back toward target. Assume $\kappa > \rho$, which makes the rule passive: real debt $\phi_t B_t^T$ is stationary under (49) for any path of monetary policy.[32] Given the fiscal rule (49), the fiscal budget constraint (5)

[32]Substituting the rule (49) into the fiscal budget constraint and using $\dot{A}_t^{BT} = r_t^a A_t^{b,T} + \phi_t \dot{B}_t^T$, real Treasury debt $A_t^{b,T} \equiv \phi_t B_t^T$ obeys $\dot{A}_t^{BT} = \bar{T} + \kappa\bar{b} - R_t - (\kappa - r_t^b) A_t^{b,T}$, which mean-reverts (toward $A_t^{b,T} = (\bar{T} + \kappa\bar{b} - R_t)/(\kappa - r_t^b)$) when $\kappa > r_t^b$ and explodes when $\kappa < r_t^b$. Since in equilibrium $r_t^b < \rho$, the condition $\kappa > \rho$ is sufficient to guarantee $\kappa > r_t^b$, hence stationarity, along any monetary path. Fiscal solvency is then independent of the central bank, which is what lets us study its problem in isolation. The implied stationary transfer is $T_{ss} = -r_{ss}^b A_{ss}^{b,T}$, the surplus that services the debt.

determines the path of Treasury debt B_t^T:

$$\phi_t \dot{B}_t^T = \bar{T} + \kappa \bar{b} - R_t - (\kappa - i_t^b)\, \phi_t B_t^T. \tag{50}$$

Second, I assume that the central bank makes no remittances, $R_t = 0$, so its seigniorage and interest income can fund only open-market purchases of Treasury debt, never direct transfers to the Treasury.[33] This makes the balance sheet a binding constraint and the separated problem genuinely distinct from the consolidated one: the central bank operates only through its balance sheet—money growth, which sets inflation, and in turn determines its holdings of Treasury debt B_t^{CB}, which fix the public's bond supply $B_t = B_t^T - B_t^{CB}$ and hence the liquidity premium i_t^b through bond-market clearing, which becomes:

$$\int (\omega - y_t)\, dG_t = \phi_t (B_t^T - B_t^{CB}). \tag{51}$$

Later, I relax this constraint to allow for $\underline{R} \leq R_t \leq \bar{R}$, where $\underline{R} \leq 0 \leq \bar{R}$. The upper bound $\bar{R}$ caps the remittances the central bank can send to the Treasury. The lower bound $\underline{R}$ caps the recapitalization it can receive in return.

4.1 Central bank's problem

The planner's —central bank's— policy choices now reduce to the size and composition of its balance sheet, $\mathbf{P} = \{\dot{M}_t, \dot{B}_t^{CB}\}_t$. Its problem is the analog of (22)–(25), extended to incorporate the behavior of fiscal policy. The saddle-point problem is

$$\mathcal{W}_0(F_0, M_0, B_0^T, B_0^{CB}) = \max_{\mathbf{X},\mathbf{P}} \min_{\mathbf{M}} \mathbb{E}_0 \left[L(0, \tau_1, F_0, M_0, B_0^T, B_0^{CB}) + e^{-\rho \tau_1} \mathcal{W}_{\tau_1}(F_{\tau_1}, M_{\tau_1}, B_{\tau_1}^T, B_{\tau_1}^{CB}) \right], \tag{52}$$

[33] Were the central bank free to choose R_t, the separated problem would collapse to the consolidated one. Summing (5)–(6) eliminates R_t and returns the consolidated budget (7), with the added fiscal rule, which the central bank can undo with its balance sheet. Holding Treasury debt $A_{ss}^{b,CB} = \bar{b} - (T^* - \bar{T})/\kappa - A_{ss}^{b*}$ makes the rule (49) return exactly T^*, where (A_{ss}^{b*}, T^*) are the consolidated optimum's public-bond supply and transfer. The position may be negative, in which case the central bank issues bonds rathe r than holds them.

with Lagrangian:

$$
\begin{aligned}
L(t_1, t_2, F_t, M_t, B_t^T, B_t^{CB}) = \int_{t_1}^{t_2} e^{-\rho s} \Bigg\{ & \int \mathcal{U}_s(x) dG_s(x) \\
& + \Big\{ \psi_s(x) \Big[-\rho W_s(x) + \mathcal{U}_s(x) + \partial_s W_s(x) + \mathcal{A}_s(x) W_s(x) \Big] \\
& + \chi_s(x) \Big[u'(c_s) - \partial_\omega W_s(x) \Big] \\
& + \eta_s(x) \Big[\alpha U'(y_s) - \alpha\, \partial_\omega W_s(\omega - y_s, z) - i_s^b\, \partial_\omega W_s(\omega, z) \Big] \mathbb{1}\{\omega > \underline{\omega}_s(z)\} \\
& + \gamma_s(x) \Big[\mathcal{A}_s^*(x)[G_s(x)] - \partial_s G_s(x) \Big] \Big\} dx \\
& + \zeta_s \Big[\phi_s \dot{B}_s^T - \bar{T} - \kappa \left(\bar{b} - \phi_s B_s^T \right) - i_s^b \phi_s B_s^T \Big] \\
& + \vartheta_s \Big[\dot{M}_s \phi_s + i_s^b \phi_s B_s^{CB} - \phi_s \dot{B}_s^{CB} \Big] \\
& + \mu_s \Big[\int y_s\, dG_s(x) - \phi_s M_s \Big] \\
& + \varphi_s \Big[\int (\omega - y_s)\, dG_s(x) - \phi_s (B_s^T - B_s^{CB}) \Big] \Bigg\} ds.
\end{aligned} \tag{53}
$$

The controls are $\mathbf{X} = \{c_t(x), y_t(x), W_t(x), G_t(x), \phi_t, i_t^b, \dot{B}_t^T\}_{t=0}^{\tau_1}$, and the multipliers are $\mathbf{M} = \{\psi_t(x), \chi_t(x), \eta_t(x), \gamma_t(x), \zeta_t, \vartheta_t, \mu_t, \varphi_t\}_{t=0}^{\tau_1}$. The fiscal rule (49) is also substituted into the household budget constraints, so the generators $\mathcal{A}_s$ and $\mathcal{A}_s^*$ depend on the aggregates ϕ_s and B_s^T directly through their effect on transfers $T_s = \bar{T} + \kappa(\bar{b} - \phi_s B_s^T)$. The household block of the Lagrangian—the HJB (multiplier ψ_s), the two Euler implementability constraints (χ_s, η_s), and the KFE (γ_s)—is identical to the consolidated problem (25); the separation changes only the government side: the single government budget splits into a Treasury budget (with multiplier ζ_s) and the central bank's budget (with multiplier ϑ_s), and bond-market clearing now references bonds in the hands of the public, $B_s = B_s^T - B_s^{CB}$, rather than total bonds outstanding. The real values of the asset stocks are $A_s^a \equiv \phi_s M_s$ (money), $A_s^{b,T} \equiv \phi_s B_s^T$ (Treasury debt), and $A_s^{b,CB} \equiv \phi_s B_s^{CB}$ (central-bank holdings); real bonds in the hands of the public are then $A_s^b \equiv \phi_s B_s = A_s^{b,T} - A_s^{b,CB}$, and total real public liabilities are $A_s \equiv A_s^a + A_s^b = A_s^a + A_s^{b,T} - A_s^{b,CB}$.

4.2 The optimality system

The first-order conditions of the Lagrangian (53) follow closely those of the consolidated problem. Proposition 5 states the system in full, mirroring Proposition 1. The discussion below focuses on the differences; the detailed derivations are in Appendix D.

Proposition 5 (Separated optimality system)**.** *For $\lambda \in \{0, \infty\}$, the optimal controls, policy instruments, and multipliers* $(\mathbf{X}, \mathbf{P}, \mathbf{M})$ *satisfy the equilibrium conditions of Definition 1 and the fiscal rule (49) together with:*

i) ***Social value.*** *γ_s solves the backward HJB*

$$\rho\gamma_s = \mathcal{U}_s - \zeta_s T_s + \mu_s y_s + \varphi_s(\omega - y_s) + \mathcal{A}_s[\gamma_s] + \partial_t \gamma_s, \tag{54}$$

with terminal condition at the handoff date $\gamma_{\tau_1}(x) = \partial \mathcal{W}_{\tau_1}/\partial F_{\tau_1}(\tilde{\omega}, z)$, where $\tilde{\omega} = \omega/\phi_{\tau_1}$.

ii) ***Revenue shadow prices.*** *The central bank's shadow price ϑ_s satisfies*

$$\dot{\vartheta}_s = (\rho - r^a_s)\,\vartheta_s - \mu_s, \tag{55}$$

$$\dot{\vartheta}_s = (\rho - r^a_s - i^b_s)\,\vartheta_s - \varphi_s, \tag{56}$$

and the Treasury's shadow price ζ_s satisfies

$$\dot{\zeta}_s = (\rho - r^b_s)\,\zeta_s - \varphi_s + \kappa(\zeta_s - \theta_s), \qquad \theta_s \equiv \int \left(\partial_\omega \gamma_s + \frac{\psi_s}{g_s}\, u'(c_s)\right) dG_s, \tag{57}$$

with terminal conditions at the handoff date $\vartheta_{\tau_1} = -(1/\phi_{\tau_1})\,\partial \mathcal{W}_{\tau_1}/\partial M_{\tau_1}$ and $\zeta_{\tau_1} = -(1/\phi_{\tau_1})\,\partial \mathcal{W}_{\tau_1}/\partial B^T_{\tau_1}$.

iii) ***Household optimality.*** *(χ_s, η_s) satisfy (31)–(32), with the central bank's shadow price ϑ_s in place of θ_s in the portfolio component of $\tilde{\eta}_s$.*

iv) ***Promise multiplier.*** *Under discretion, $\psi_s \equiv 0$. Under commitment, ψ_s solves the forward KFE (33) with initial condition $\psi_0(x) = 0$, unchanged from the consolidated system.*

v) ***Optimal prices.*** *The liquidity premium and the price of money satisfy*

$$\int b_s(\partial_\omega \gamma_s\, g_s + \psi_s\, u'(c_s))\, dx \;=\; \zeta_s A^{b,T}_s - \vartheta_s A^{b,CB}_s \;+\; \int \eta_s\, u'(c_s)\, dx, \tag{58}$$

$$\int \omega\,(\partial_\omega \gamma_s\, g_s + \psi_s\, u'(c_s))\, dx \;=\; \vartheta_s A^a_s + \zeta_s A^{b,T}_s - \vartheta_s A^{b,CB}_s, \tag{59}$$

where $b_s = \omega - y_s$.

The system is nearly identical to the consolidated optimality system of Proposition 1, with the exception that under separation the planner accounts for which government budget constraint revenue and spending show up in, with different shadow values. Seigniorage accrues to the central bank, valued according to ϑ_s. Bond interest payments leave the Treasury budget but arrive in the Central Bank budget. Transfers are paid by the Treasury, so ζ_s prices them.

The scalar θ_s in (57) retains its consolidated interpretation as the social value of a uniform transfer, but under separation it equals neither budget multiplier. With no free transfer to op-

timize over, the first-order condition that would tie the budget multipliers to θ_s is absent. The fiscal-rule feedback $\kappa(\zeta_s - \theta_s)$ prices the rule. As debt rises, transfers fall, which eases the Treasury budget, valued at ζ_s, and cuts household wealth, valued at θ_s.

Part v) contains the change that is most consequential for the results that follow. In the liquidity-premium condition (58), the fiscal cost of a higher premium is the net inter-budget interest cost $\zeta_s A_s^{b,T} - \vartheta_s A_s^{b,CB}$, the Treasury's interest expense net of the central bank's interest income. In the optimal-inflation condition (59), the devaluation revenue accrues separately, the central bank gaining on its net money issuance and the Treasury on the dilution of its real debt. Both conditions collapse to their consolidated counterparts when $\zeta_s = \vartheta_s$. The wedges among the three shadow prices are the objects that organize the next two subsections.

The remaining parts carry over. The promise multiplier and the implementability multipliers are unchanged in form, and the system retains its forward–backward structure, with λ entering only through the conditions on ψ_s. At the handoff date each backward co-state matches the successor's valuation of the state it prices, and because the successor cannot re-mix liabilities across the two budgets, ϑ_s and ζ_s are generically distinct. Appendix D contains the derivations and boundary conditions.

4.3 Optimal central bank balance sheets with full discretion

Under discretion, the central bank at each time period is constrained only by the state $(F_t, M_t, B_t^T, B_t^{CB})$ it inherits, not by past promises ($\psi_s \equiv 0$). In a stationary equilibrium: real aggregates are constant while the nominal stocks grow at the common rate $\dot{M}_{ss}/M_{ss} = \dot{B}_{ss}^T/B_{ss}^T = \dot{B}_{ss}^{CB}/B_{ss}^{CB} = -r_{ss}^a$.

With no remittance, $R = 0$, the central bank is constrained to set policy through open-market operations—swapping money for Treasury bonds on its balance sheet—so its balance-sheet share B^{CB}/B^T fixes the bonds in the hands of the public and directly affects the liquidity premium i^b through bond-market clearing. The real return on money, $-r^a$, however, is not a separate instrument: the central bank can only affect it through its balance sheet, which must be self-financing. In steady state, its budget constraint equates the real return r_{ss}^a paid on its money liabilities A_{ss}^a to the real bond return $r_{ss}^b \equiv r_{ss}^a + i_{ss}^b$ earned on its asset holdings $A_{ss}^{b,CB}$,

$$r_{ss}^a \, A_{ss}^a = r_{ss}^b \, A_{ss}^{b,CB}. \tag{60}$$

This is the central distinction from the consolidated problem. Under separation, the budget constraint (60) leaves the central bank with effectively a single instrument: choose the money growth rate to achieve a stationary inflation rate, then (60) then pins the nominal interest rate i_{ss}^b through the effect on debt held in the hands of the public. Or, vice-versa, target a liquidity premium and the separated central bank gives up an independent inflation target. The central bank must therefore trade off the inflation tax against the liquidity premium along the budget

line.

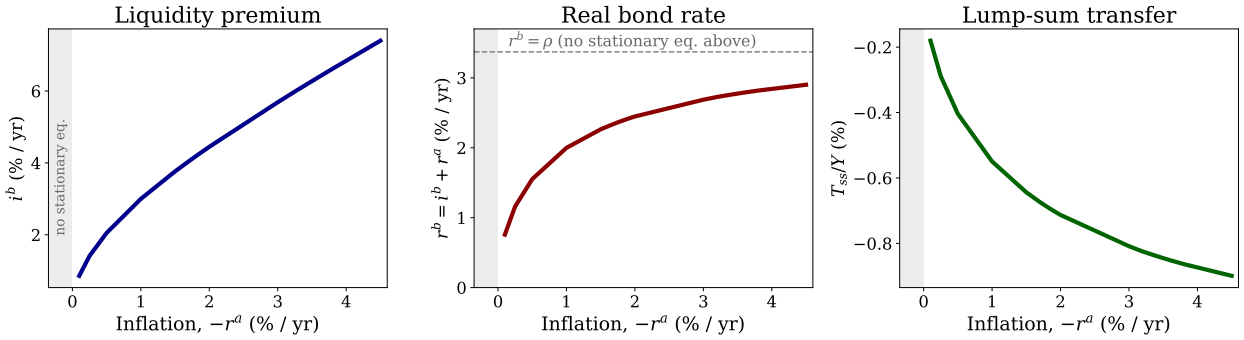


Figure 9: The separated equilibrium locus.

Figure 9 illustrates this trade-off under the calibration from Section 2.5. The Treasury's real debt target is set to match the calibration target of 22% of GDP, $\bar{b} = 0.22 \times Y$. The parameter κ is the Treasury's reaction of the primary surplus to real debt above target, which Bohn (1998) estimates between 0.028 and 0.054, depending on the sample period within 1916–1995, and finds rising with the level of debt. I use $\kappa = 0.10$.[34] The left panel shows the liquidity premium rises sharply with inflation, from 0.9 percent at around price stability to 7.4 percent at four and a half percent inflation. The middle panel plots the real bond return along the same locus. Since the liquidity premium rises more than one-for-one as inflation rises (r^a_{ss} falls), the real bond return rises with inflation. The right panel illustrates the fiscal effect on transfers.

How does the discretionary central bank weigh the trade-off between inflation and the liquidity premium? The central bank's optimality conditions for the liquidity premium and inflation, (58) and (59) respectively, are now not independent since they share a common inter-budget cost. Writing these conditions in covariance form—subtracting the average social value θ_{ss} on each household asset position—and differencing, gives the separated analog of the consolidated discretion optimum of Proposition 3.

Proposition 6 (Optimal separated policy under full discretion)**.** *Under Assumptions 1 and 2, full discretion, and the passive fiscal rule* (49)*, a stationary optimum satisfies*

$$\underbrace{(\vartheta_{ss} - \theta_{ss})\, A^a_{ss} - \mathrm{Cov}_{G_{ss}}(\omega,\, \partial_\omega \gamma_{ss})}_{\textit{inflation margin}} \;=\; \underbrace{\int \eta_{ss}\, u'(c_{ss})\, dx - \mathrm{Cov}_{G_{ss}}(\omega - y_{ss},\, \partial_\omega \gamma_{ss})}_{\textit{premium margin}}, \tag{61}$$

where $\mathrm{Cov}_{G_{ss}}$ *is the cross-sectional covariance under* G_{ss}*,* $\theta_{ss} = \int \partial_\omega \gamma_{ss}\, dG_{ss}$ *is the social value of a uniform transfer, and* η_{ss} *is the lumpy implementability multiplier* (32)*.*

Marginally higher inflation leads to a higher liquidity premium, along the locus of Figure 9. The discretionary central bank understands that any redistributive gain from inflation, equal to

[34] D'Erasmo et al. (2016) estimate $\kappa = 0.078$ with data through 2014.

$-\text{Cov}_{G_{ss}}(\omega, \partial_\omega \gamma_{ss})$, must be evaluated against the redistributive cost of higher interest payments to wealthy bond owners, equal to $-\text{Cov}_{G_{ss}}(\omega - y_{ss}, \partial_\omega \gamma_{ss})$. The two redistributive incidences that the consolidated planner can target separately, in Proposition 3, now work against each other and net to $\text{Cov}_{G_{ss}}(y_{ss}, \partial_\omega \gamma_{ss})$, the cross-sectional incidence on money holdings alone. Since the distribution of money holdings is more condensed than for bonds or total wealth, this redistributive effect is naturally much smaller. In addition, a higher premium carries the liquidity cost $\int \eta_{ss} u'(c_{ss})\, dx$, as before, and the additional seigniorage from higher inflation is evaluated at the wedge $(\vartheta_{ss} - \theta_{ss})$ on the money base. This last term is new and captures that if the central bank's value of seigniorage is below the social value, higher inflation is costly.

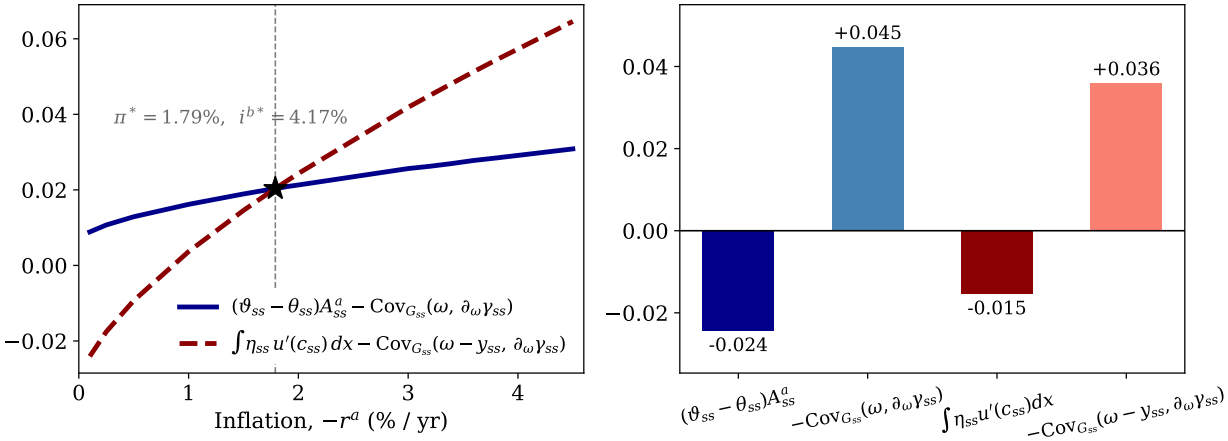


Figure 10: The combined optimality condition under separated discretion.

Figure 10 illustrates these margins under the calibration of Section 2.5. The left panel shows a discretionary optimum at $\pi^* = 1.79$ percent, where the market-clearing premium is $i^{b*} = 4.17$ percent. To the left of the equilibrium, the marginal benefit of higher inflation exceeds the marginal cost of a higher liquidity premium and the central bank gains from faster money growth. To the right, the opposite is true. The right panel reports the four terms of (61) at the optimum. The two covariance terms dominate in magnitude but largely cancel, +0.045 against +0.036, because their bases differ only by money holdings. The net marginal redistributive benefit of higher inflation is the difference $-\text{Cov}_{G_{ss}}(y_{ss}, \partial_\omega \gamma_{ss}) = 0.009$. What keeps inflation in check for the discretionary central bank is the net of the fiscal wedge, $(\vartheta_{ss} - \theta_{ss})A^a_{ss} = -0.024$, and the liquidity cost of a higher premium, $\int \eta_{ss} u'(c_{ss})\, dx = -0.015$.

The fiscal wedge $\vartheta_{ss} - \theta_{ss}$ measures the difference between the value of real seigniorage inside the central bank's budget and the value of seigniorage transferred uniformly to households. To sign it, first solve for the social value of wealth from the central bank's standpoint by substituting the steady-state transfer $T_{ss} = -(r^a_{ss} + i^b_{ss})A^{b,T}_{ss}$ and the co-states (55)–(56) into (54), and using the

budget (60):

$$\rho\gamma_{ss} = \mathcal{U}_{ss} + \rho\vartheta_{ss}A_{ss} + \underbrace{(\rho - r^a_{ss})\vartheta_{ss}[\omega - A_{ss}]}_{\text{inflation-tax redistribution}} - \underbrace{i^b_{ss}\,\vartheta_{ss}[(\omega - y_{ss}) - A^b_{ss}]}_{\text{liquidity-premium redistribution}} + \underbrace{r^b_{ss}A^{b,T}_{ss}\,(\zeta_{ss} - \vartheta_{ss})}_{\text{inter-budget revaluation}} + \mathcal{A}_{ss}[\gamma_{ss}]. \tag{62}$$

The social value of wealth mirrors that of the consolidated planner in (42), with two differences. The redistributive terms are now valued at the central bank's shadow price ϑ_{ss} rather than θ_{ss}, and a new revaluation term.[35] Notice, every term in (62) that varies with individual wealth is priced at the central bank's shadow value ϑ_{ss}. The argument of Lemma 5 therefore applies with ϑ_{ss} in place of θ_{ss}:

$$\partial_\omega\gamma_{ss} = u'(c_{ss}) + \vartheta_{ss}\,\partial_\omega\gamma^f_{ss}, \tag{63}$$

with the same fiscal externality γ^f_{ss} as in the consolidated problem. With no remittances, the user-cost revenue a household's wealth generates accrues to the central bank, so it is the central bank's valuation that prices it. Averaging (63) across households gives $\theta_{ss} = \int u'(c_{ss})\,dG_{ss} + \vartheta_{ss}\bar{F}_{ss}$, where $\bar{F}_{ss} = \int \partial_\omega\gamma^f_{ss}\,dG_{ss}$ is the average marginal fiscal value of household wealth, and therefore

$$\vartheta_{ss} - \theta_{ss} = \vartheta_{ss}(1 - \bar{F}_{ss}) - \int u'(c_{ss})\,dG_{ss}. \tag{64}$$

The wedge compares the two uses of a public dollar that separation severs. A dollar transferred to households buys consumption at average marginal utility and returns the fraction $\bar{F}_{ss}$ of itself as future user-cost revenue. A dollar inside the central bank buys only the revenue services. The wedge is negative whenever the central bank's dollar is worth less than average marginal utility; $\vartheta_{ss} \leq \int u'(c_{ss})\,dG_{ss}$ is sufficient.

What determines the sign of the wedge is the stance of fiscal policy. Were remittances unconstrained, the wedge would close: the choice of R_t would equate ζ_{ss} and ϑ_{ss}, and the fiscal rule's feedback would pass the equality on to θ_{ss}, so a dollar inside the central bank could always become a dollar in households' hands. The zero-remittance statute severs that conversion, and the wedge records which side the rule leaves short. A generous rule leaves households wealthy and average marginal utility low. Government revenue is then the scarcer object, and by (64) the wedge tends positive, rewarding inflation. The calibrated rule is the opposite case. In a stationary equilibrium it collects a net tax, $T_{ss} = -r^b_{ss}A^{b,T}_{ss} < 0$, so households are the short side and average marginal utility is high. The wedge is negative and acts to constrain inflation.

[35] Each household's tax payment $-T_{ss}$ services Treasury debt, a fiscal benefit valued at ζ_{ss}, while the central bank values resources at ϑ_{ss}, so the payment carries the revaluation $(\zeta_{ss} - \vartheta_{ss})$.

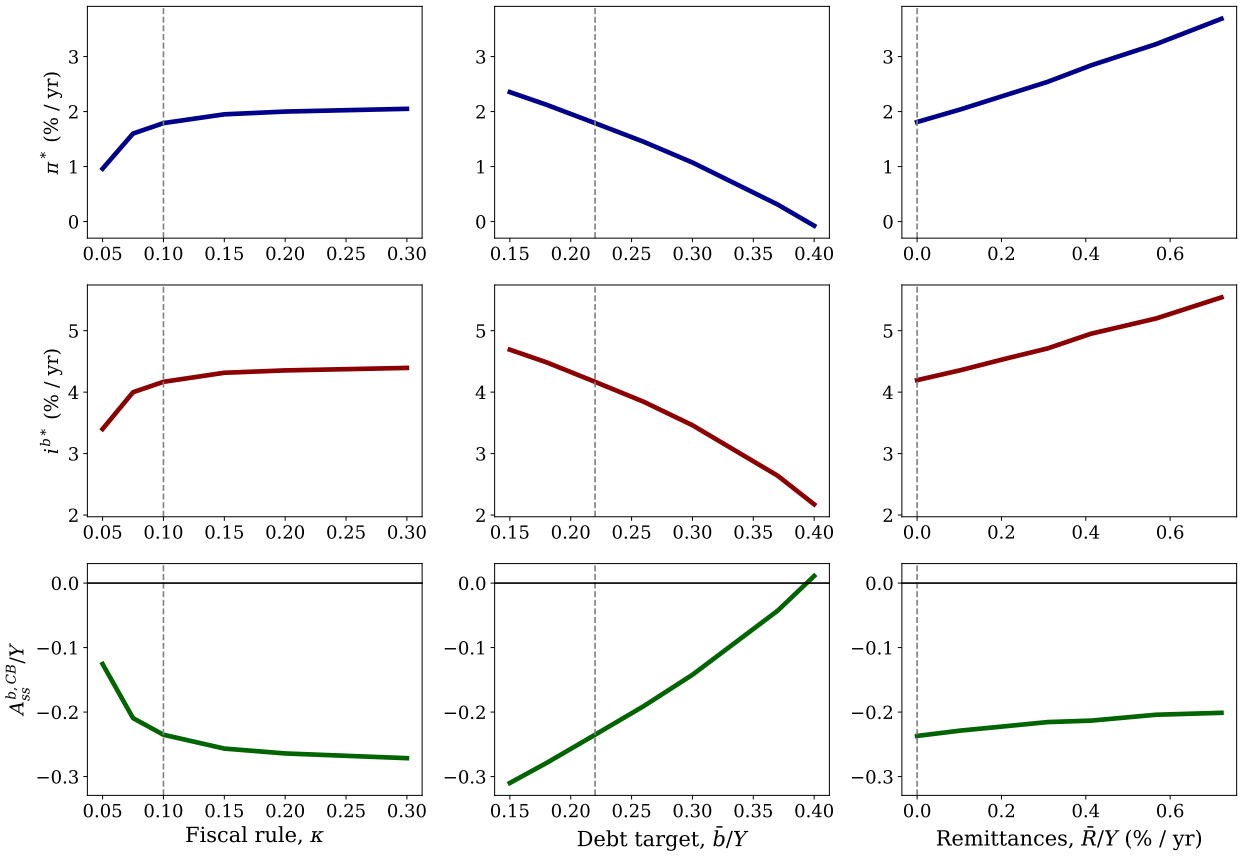


Figure 11: Comparative statics of the separated stationary optimum.

Comparative statics. Figure 11 shows comparative statics of the discretionary stationary optimum (π, i^b) and the size of the central bank's balance sheet to GDP, $A^{b,CB}/Y$, with respect to fiscal policy. The left column varies the fiscal reaction κ. Across reactions from 0.05 to 0.30, optimal inflation rises by about a percentage point, from 0.96 to 2.05 percent, flattening above $\kappa \approx 0.15$. A weak fiscal reaction lets debt drift above target, and the larger bond supply pulls optimal inflation down. The debt target matters more (middle column). Scarcer debt raises optimal inflation, the premium, and the size of the short position. When the Treasury supplies less of the bond stock households want to hold, the central bank supplies the remainder and services it with a larger inflation levy. In the other direction, at a debt target near forty percent of GDP the optimum reaches price stability, $\pi^* = -0.1$ percent, with an essentially empty balance sheet, and no interior optimum exists beyond it. The right column raises the limit on remittances the central bank can transfer to the Treasury $R_t \leq \bar{R}$. Optimal inflation rises with $\bar{R}$ since the central bank wants $R = \bar{R}$, and the short position shrinks. The stationary point disappears near $\bar{R} \approx 0.7$ percent of GDP, before the central banks' position turns positive.

The sign of the balance sheet. Across all three experiments the central bank's stationary bond position is short, which is counterfactual. The sign is an artifact of a stationary optimum featuring $r^a_{ss} < 0 < r^b_{ss}$. With zero remittances, stationarity forces the interest on the bond position to absorb the seigniorage flow, $r^b_{ss} A^{b,CB}_{ss} = r^a_{ss} A^a_{ss}$, so any inflationary stationary policy with a positive real bond return requires $A^{b,CB}_{ss} < 0$. The computed locus delivers $r^b_{ss} > 0$ throughout, so the condition is not restrictive here. One way to generate a long central bank bond position is to allow for growth. In a growing economy the same accounting delivers the opposite sign. Keeping real balances in line with output requires nominal money growth $\pi + g$, and the seigniorage this generates finances a growing asset portfolio rather than interest on a short position: the stationary position becomes $A^{b,CB} = (\bar{R} - (\pi + g)A^a)/(r^b - g)$, which is positive for any remittance below seigniorage whenever $r^b < g$—the empirically relevant case.[36] The short positions here are the zero-growth image of the actual arrangement, in which central banks remit profits and hold assets that grow with money demand.

4.4 Optimal central bank balance sheet with full commitment

Commitment changes the separated problem exactly as it changes the consolidated one (Section 3.5): the promise multiplier ψ_{ss} is no longer zero, and the same wealth- and bond-weighted promise channels—the commitment premiums of Proposition 4—enter the optimal conditions, attached to the redistributive (covariance) side. The inter-budget wedges are untouched. Writing the two price conditions in covariance form and differencing, exactly as in the discretion case, gives the combined stationary condition.

Proposition 7 (Optimal separated policy under full commitment)**.** *Under Assumptions 1 and 2, full commitment, and the passive fiscal rule* (49)*, a stationary optimum satisfies*

$$\underbrace{(\vartheta_{ss} - \theta_{ss})\, A^a_{ss} - \mathrm{Cov}_{G_{ss}}(\omega,\, \partial_\omega \gamma_{ss}) - \Psi_{ss}}_{\textit{inflation margin}} \;=\; \underbrace{\int \eta_{ss}\, u'(c_{ss})\, dx - \mathrm{Cov}_{G_{ss}}(\omega - y_{ss},\, \partial_\omega \gamma_{ss}) - \Psi^b_{ss}}_{\textit{premium margin}}, \tag{65}$$

where the commitment premiums $\Psi_{ss} = \int (\omega - A_{ss})\, \psi_{ss}\, u'(c_{ss})\, dx$ *for total wealth and* $\Psi^b_{ss} = \int ((\omega - y_{ss}) - A^b_{ss})\, \psi_{ss}\, u'(c_{ss})\, dx$ *for illiquid wealth are those of Proposition 4,* θ_{ss} *is the promise-augmented social value of a uniform transfer from* (57)*, and* ψ_{ss} *solves the forward KFE* (33) *with* $\partial_t \psi_{ss} = 0$*.*

Condition (65) adds a commitment premium subtracted from each margin. The premiums are exactly the consolidated ones, so their interpretation in Section 3.5 carries over unchanged. In the combined condition (65), the two premiums net to $\Psi_{ss} - \Psi^b_{ss} = \int (y_{ss} - A^a_{ss})\, \psi_{ss}\, u'(c_{ss})\, dx$,

[36] In 2023 the Federal Reserve's securities portfolio earned $163.8 billion on average holdings near $7.9 trillion (Board of Governors of the Federal Reserve System, 2024), a nominal yield of about 2.1 percent against nominal output growth of 6.3 percent (U.S. Bureau of Economic Analysis, 2024). The portfolio's real return averaged roughly 0.5–1.5 percent in the decades before 2021 and turned negative during the inflation of 2021–2023.

the promise content carried on money holdings—the promise analog of the covariance netting in Proposition 6. The term is positive at every stationary equilibrium, so honoring promises tilts the trade-off toward less inflation and the lower premium that comes with it. The resulting stationary optimum is at a corner where the premium reaches zero and $\pi = -0.5\%$.

4.5 Summary: optimal stationary policy in the four regimes

Table 3 collects the stationary optima of the four regimes of Sections 3 and 4. The tension a planner faces between raising fiscal revenue to provide public insurance and preserving the returns households respond to when they save to self-insure, is resolved in different directions depending on the ability to commit. Under consolidated discretion, the tension resolves entirely toward inflationary redistribution so much so the economy heads to demonetization. Under consolidated commitment, the tension resolves in the other direction pushing returns as high as possible, stopping just short of the Friedman rule. Separating the central bank from the Treasury provides a brake in either direction. Separation forces a discretionary planner to trade off the redistributive benefits of inflation with the redistributive costs of a higher liquidity premium. The result is a moderate stationary inflation rate and positive liquidity premium. For a committed planner, separation halts the push toward the Friedman rule at the zero-premium corner, where the premium hits its floor before inflation close to Friedman deflation.

	Discretion	Commitment
Consolidated	none—inflates toward demonetization	$\pi^* = -3.35\%$, $i^{b*} = 0$
Separated	$\pi^* = 1.79\%$, $i^{b*} = 4.17\%$, $M/B \approx 0.59$	corner: $\pi^* = -0.5\%$, $i^{b*} = 0$, $M/B \approx 1.3$

Calibration for comparison: $\pi = 2\%$, $i^b = 2\%$, $M/B = 1.37$.

Table 3: Stationary optima across the four regimes.

5 Optimal time-consistent central bank responses to shocks

This section studies how the separated discretionary central bank optimally responds to two types of unanticipated shocks: (i) a negative total factor productivity shock and (ii) a temporary increase in unemployment risk. Together, these shocks highlight two channels through which aggregate shocks reach the planner's problem. The productivity contraction changes total resources available in the economy by scaling all incomes proportionately. The rise in unemployment risk also reduces resources but also introduces a direct redistributive effect; job loss becomes both more likely and longer lasting, which concentrates losses on the unemployed and strengthens the precautionary motive of everyone else. For both, I compute perfect-foresight transitions

around the stationary optimum of Section 4.3. The transitions solve the time-varying optimality system of Appendix E.3 by the sequence-space Jacobian method of Appendix E.4.[37]

5.1 A total factor productivity contraction

Figure 12 shows the optimal response to an unanticipated one percent fall in total factor productivity that decays exponentially, the remaining gap proportional to e^{-t} with t in years, so roughly a third of the contraction remains after one year. The solid blue lines illustrate the optimal dynamic responses. The dotted and dashed lines illustrate the equilibrium response if instead the central bank fixed one instrument at its stationary value, the inflation rate and the liquidity premium, respectively. Since the central bank has only one policy tool through its balance sheet, fixing inflation lets the liquidity premium be determined in equilibrium by market clearing. The same is true fixing the liquidity premium.

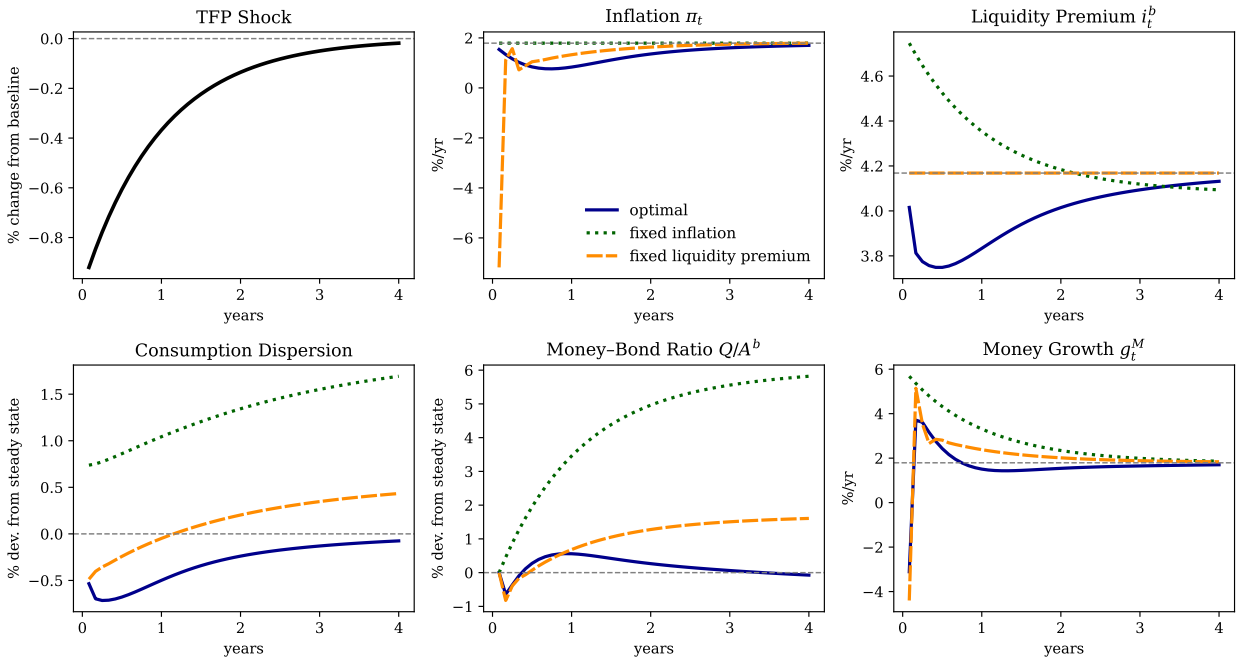


Figure 12: Impulse response to a one percent productivity contraction.

The central bank meets the productivity contraction with disinflation, π_t falling from 1.79 to 1.54 percent on impact and hitting a minimum at 0.76 percent around nine months, before increasing back toward its stationary level. A negative TFP shock reduces the nominal wealth base that the discretionary central bank can tax through inflation, while the redistributive benefit of the temptation barely moves, so the optimal inflation levy falls. The contraction scales all incomes proportionately, so the shape of the wealth distribution is roughly unchanged, and

[37] The transition system is linearized in the size of the shock. Appendix E.4 gives details.

with constant relative risk aversion a proportional fall in consumption raises every household's marginal valuation by a common factor. The bottom-left panel illustrates this; consumption dispersion falls roughly three quarters as much as the contraction and recovers with it.

The optimal liquidity premium also falls, from 4.17 to a trough of 3.75 percent, and recovers slowly. The premium is the levy on liquidity, and the same logic applies. Lower incomes shrink the demand for liquid balances, so the base of the levy falls, while the proportional contraction leaves its redistributive benefit per unit of base unchanged. Because marginal valuations are elevated during the contraction, the utility cost of the premium is temporarily high, and the central bank cheapens liquidity when it is most valuable. The bottom middle and right panels show how the optimal policy maps into the instruments, the money-bond ratio and money growth, respectively. Under the optimal policy the composition barely moves, while both fixed policies swing it several percent away from its stationary level.

Both responses also have a dynamic component. The social valuations that price the two levies are forward looking, and wealth carried through the contraction has elevated social value because it buffers the remaining downturn and rebuilds the tax base for the recovery. The planner therefore keeps both levies low while wealth is depleted, which is why the responses hit their minimums months after the shock rather than at impact.

Suppose instead the central bank fixes a policy instrument. Fixed inflation leads to a rise in the liquidity premium since the fall in TFP reduces real wealth (including bond) demand but the supply adjusts slowly via the fiscal rule. With the return on money pinned by the inflation target, the bond return must rise to keep the outstanding stock willingly held. The premium climbs six tenths of a percentage point above its stationary level and real balances fall just over four percent, more than twice the optimal decline. Fixing the liquidity premium instead removes the bond-return margin, and the whole adjustment loads onto inflation. The disinflation is front loaded, arriving as a one-time fall in the price level of about three quarters of a percent, and recovers faster. In both cases the burden of the shock is relocated rather than reduced. The balance sheet composition drifts far from its stationary level, while the optimal policy spreads the adjustment across both prices and keeps both distortions small. Under either fixed policy, consumption dispersion is higher than is socially optimal (bottom-left panel).

5.2 A rise in unemployment risk

This section considers a shock to unemployment risk (a fall in the job-finding rate and an increase in the separation rate) that raises unemployment from six to near ten percent. In line with the results of Shimer (2005, 2012), I assume two thirds of the movement in unemployment is generated by the finding rate while one third is through the separation rate.[38] Figure 13 shows

[38] At impact the job-finding hazard falls from 4.28 to 1.85 per year, a decline of 57 percent, and the separation hazard rises from 0.27 to 0.42 per year, an increase of 52 percent. Both revert with the same e^{-t} decay as the productivity

the responses in the same format as Figure 12.

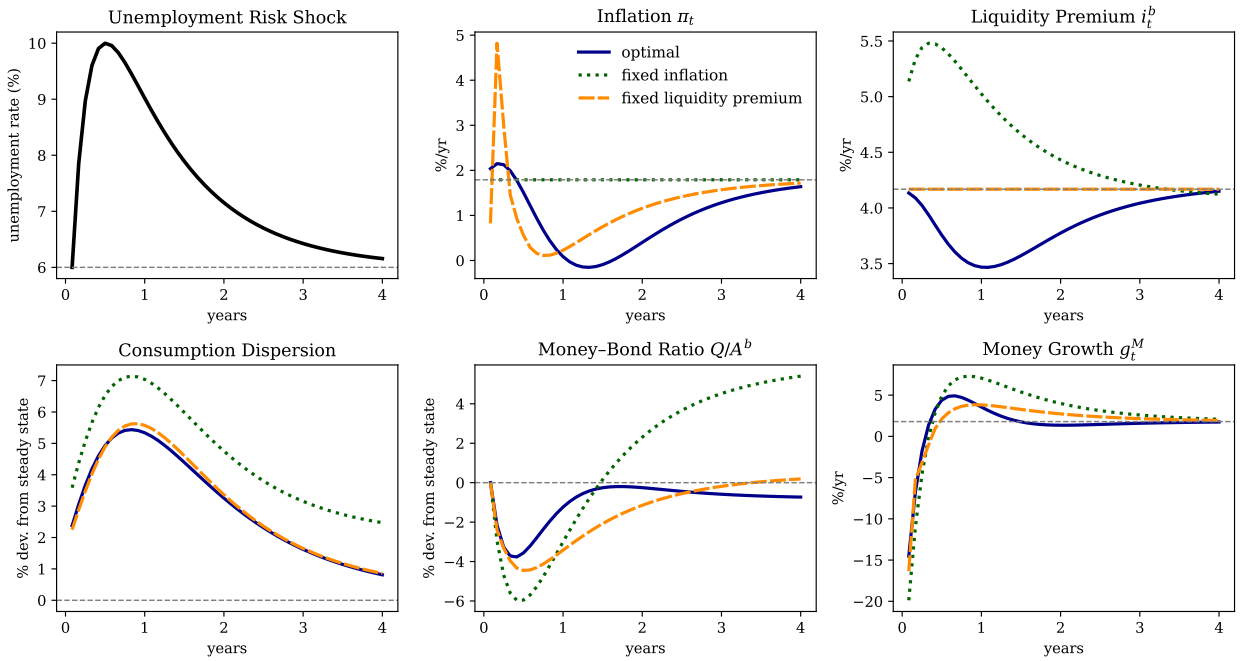


Figure 13: Impulse response to a rise in unemployment risk. Unemployment peaks at ten percent five months after impact.

The central bank responds to the rise in unemployment risk by providing insurance through lower inflation and liquidity premia. Inflation rises about a quarter of a percentage point at impact and then falls steeply, turning into deflation and reaching roughly -0.2% before recovering. The liquidity premium falls from 4.17 to 3.47 percent over the same horizon.

While the responses look similar, the two levies fall for a different reason than under the TFP contraction. Here, the redistributive incentive to inflate rises, because losses concentrate on the unemployed and the social value of wealth steepens at the bottom of the distribution. However, the same shock raises every household's demand for a buffer—since job loss is now more likely and spells last longer. In response, the central bank raises the return on money and cheapens liquidity, letting households build the buffers they demand. The initial rise in inflation is the one concession to the levy. It taxes nominal wealth once, at the moment the redistributive value of doing so is highest. The composition of the balance sheet dips as bond demand recovers first and returns to its stationary level within two years.

Fixing a policy instrument again relocates the burden. With inflation pinned, the surge in wealth demand meets a slowly adjusting bond supply and the market imposes a higher premium, which jumps above five percent, further taxing liquidity when households value it. With the premium pinned, inflation carries the adjustment and front-runs the optimal disinflation.

shock. The finding-rate decline is somewhat larger than its Great Recession counterpart.

Consumption dispersion rises under every policy, the direct incidence of the larger unemployment pool. However, the optimal response contains the rise while holding inflation fixed raises it by half again on impact, since the premium spike falls hardest on the households least able to bear it.

6 Robustness

6.1 Nominal interest on money (reserves)

The stationary optimality conditions of Sections 3 and 4 pin down the real returns on nominal wealth. Inflation is then determined directly since money pays no interest. Suppose instead the central bank pays a nominal rate i^a_{ss} on money, similar to an interest-on-reserves policy the Federal Reserve has employed since 2008. Every stationary condition holds unchanged with $r^a_{ss} = i^a_{ss} - \pi_{ss}$ in place of $r^a_{ss} = -\pi_{ss}$. The same optimum is implemented now by any pair (i^a_{ss}, π_{ss}) that delivers the optimal real return.[39] Importantly, allowing the central bank to pay nominal interest on money balances does not relax the budget constraint or balance sheet. The aggregate supply of real money balances, the composition of money and bonds, and the central bank's position are the same at every implementation. Paying interest on reserves is not a requirement to reach a zero liquidity premium.

Further, allowing interest on money does not blunt the temptation to inflate. Whatever nominal rate accrues to money, the interest is itself a nominal claim. It sits inside household nominal wealth, and a devaluation dilutes it. Households rebalance freely, so the inherited state a discretionary planner acts on is total nominal wealth, and its ex-post margin is the price level. The temptation to redistribute through surprise inflation is unchanged, and the pull toward demonetization operates exactly as before. Interest on money changes what the government pays on nominal wealth. It cannot insulate nominal wealth from the price level.

6.2 Productive real capital

In the baseline economy all wealth is nominal and exposed to inflationary devaluation. Suppose instead household also have access to real productive capital. Firms produce with constant returns from capital and labor, capital depreciates, and households accumulate it freely alongside money and bonds. I assume bonds and capital are equally illiquid, so arbitrage ties the real bond return to the net marginal product of capital, and neither the Treasury or the central bank holds private assets. Appendix F details the environment and re-derives the optimality systems.

[39] The same is true of the separated problem. A separated central bank that pays interest on money expands the nominal money supply by the interest it pays, but its real budget is unchanged: the position identity holds with $r^a_{ss} = i^a_{ss} - \pi_{ss}$ in place of $r^a_{ss} = -\pi_{ss}$, and the set of real stationary configurations it can reach is the same.

The optimality systems keep their structure relative to the baseline, with two changes to the conditions that set inflation and the premium. First, the devaluation margin taxes smaller base since capital absorbs a portion of household wealth. However, though the nominal base that remains still covaries negatively with the social value of wealth exactly as in the baseline which continues to drive towards inflationary redistribution. Second, the real capital stock anchors the illiquid return to the marginal product. As a result inflation is passed into the liquidity premium one for one, and the composition of the balance sheet loses its grip on the premium. In the baseline only a separated central bank faced a trade-off between inflation and the premium. With capital every planner faces it. A third wedge prices the capital-accumulation margin, and under commitment the promise premium concentrates there.

The open question is then whether the consolidated discretionary planner still is driven to hyperinflation in a calibrated economy. I set the capital share to one third and depreciation to six percent per year, and carry over all preference and income parameters from the baseline. One parameter is re-fit, the arrival rate of expenditure shocks, which targets the same liquid balances of thirty percent of GDP. The implied capital-output ratio is 3.6, in line with fixed private assets, and the calibrated liquidity premium is 5.3 percent. I find that optimal consolidated discretionary policy still has no stationary optimum. Even though the base is an order of magnitude smaller than in the baseline, the marginal value of a surprise dilution remains positive at every inflation rate. Under commitment, the consolidated government sets inflation so that the return on money equals the return on illiquid assets, zero-premium corner. Separation under discretion continues to deliver an interior stationary optimum with deflation of 3.0 percent, a liquidity premium just below one percent, and a central bank position of roughly thirty percent of output. The baseline's optimum pairs moderate inflation with a short position; capital flips both signs. While the quantitative results change, the primary qualitative results—the asymmetry between commitment and discretion and the role of separation—remain intact.

6.3 Portfolio adjustment costs and a richer income process

Finally, I extend the model to incorporate flow portfolio adjustment costs and a richer income process, along the lines of Kaplan and Violante (2014) and Kaplan et al. (2018), which leads to a more realistic joint distribution of wealth and income. For instance, the wealth Gini rises from 0.32 in the baseline to 0.84, against 0.85 in the data. The model also expands the incidence of liquidity premium. Since money is held not only for lumpy purchases, the premium is priced on the deposit margin (to and from the illiquid account). While adding adjustment costs increases the liquidity premium, ceteris paribus, the conditions determining optimal policy continue to have the similar covariance forms. Appendix G details the environment and re-derives the optimality systems.

7 Conclusion

Should the central bank aim to have a smaller balance sheet? I study this question in an economy in which households face uninsurable income and expenditure risk and self-insure with two nominal assets, liquid money and illiquid bonds. The size of the balance sheet fixes the relative supply of the two assets in the hands of the public and, through it, the inflation rate and the liquidity premium. Optimal policy weighs the social insurance the taxes finance against the private self-insurance they discourage. The optimality conditions price that trade-off through covariances between asset holdings and the social value of wealth across the distribution of households. A central bank that can commit should run a large, money-heavy balance sheet that satiates the economy with liquidity, drives the premium to zero, and delivers mild deflation to reward the saving that self-insurance requires. A central bank that cannot commit inflates to tax nominal wealth, and runs a smaller balance sheet to keep liquidity scarce and expensive.

The analysis of the paper could be extended in several dimensions. Introducing aggregate risk is one since this opens a channel through which balance sheet policy (and its response to aggregate shocks) may need to trade-off redistribution and insurance against aggregate shocks, as Bigio and Sannikov (2026) highlight. Optimal policy in this setting is an open question. Second, the analysis abstracts from a banking sector. Household directly hold central bank liabilities (as if there was only a narrow banking sector). In practice, much of the central bank's balance sheet is held by banks and transmits to households through deposit and loan rates. Studying optimal balance sheet policy in this setting is also an interesting direction for future research.

Disclosure statement. The author has no financial relationships or other potential conflicts of interest to disclose.

References

Acharya, S. and Dogra, K. (2020). Understanding HANK: Insights from a PRANK. *Econometrica*, 88(3):1113–1158.

Acharya, V. V., Chauhan, R. S., Rajan, R. G., and Steffen, S. (2022). Liquidity dependence: Why shrinking central bank balance sheets is an uphill task. In *Reassessing the Effectiveness and Transmission of Monetary Policy*. Federal Reserve Bank of Kansas City. Jackson Hole Economic Policy Symposium.

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2022). Income and wealth distribution in macroeconomics: A continuous-time approach. *The Review of Economic Studies*, 89(1):45–86. Advance access February 2021.

Aguiar, M., Amador, M., and Arellano, C. (2024). Pareto improving fiscal and monetary policies: Samuelson in the New Keynesian model. Working paper, July 2024.

Aiyagari, S. R. and McGrattan, E. R. (1998). The optimum quantity of debt. *Journal of Monetary Economics*, 42(3):447–469.

Albanesi, S., Chari, V. V., and Christiano, L. J. (2003). Expectation traps and monetary policy. *The Review of Economic Studies*, 70(4):715–741.

Alvarez, F., Kehoe, P. J., and Neumeyer, P. A. (2004). The time consistency of optimal monetary and fiscal policies. *Econometrica*, 72(2):541–567.

Andolfatto, D. (2011). A note on the societal benefits of illiquid bonds. *Canadian Journal of Economics*, 44(1):133–147.

Angeletos, G.-M., Collard, F., and Dellas, H. (2023). Public debt as private liquidity: Optimal policy. *Journal of Political Economy*, 131(11):3233–3268.

Angeletos, G.-M., Lian, C., and Wolf, C. K. (2024). Deficits and inflation: HANK meets FTPL. Working paper, November 2024.

Arce, Ó., Nuño, G., Thaler, D., and Thomas, C. (2020). A large central bank balance sheet? Floor vs corridor systems in a New Keynesian environment. *Journal of Monetary Economics*, 114:350–367.

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

Azzimonti, M. and Yared, P. (2019). The optimal public and private provision of safe assets. *Journal of Monetary Economics*, 102:126–144.

Barro, R. J. and Gordon, D. B. (1983). A positive theory of monetary policy in a natural rate model. *Journal of Political Economy*, 91(4):589–610.

Bassetto, M. (2002). A game-theoretic view of the fiscal theory of the price level. *Econometrica*, 70(6):2167–2195.

Bassetto, M. and Cui, W. (2024). A ramsey theory of financial distortions. *Journal of Political Economy*, 132(8).

Benigno, P. (2020). A central bank theory of price level determination. *American Economic Journal: Macroeconomics*, 12(3):258–283.

Benigno, P. and Nisticò, S. (2017). Safe assets, liquidity, and monetary policy. *American Economic Journal: Macroeconomics*, 9(2):182–227.

Benigno, P. and Nisticò, S. (2020). Non-neutrality of open-market operations. *American Economic Journal: Macroeconomics*, 12(3):175–226.

Berentsen, A., Camera, G., and Waller, C. (2007). Money, credit and banking. *Journal of Economic Theory*, 135:171–195.

Bewley, T. F. (1980). The optimum quantity of money. In Kareken, J. H. and Wallace, N., editors, *Models of Monetary Economies*, pages 169–210. Federal Reserve Bank of Minneapolis, Minneapolis, MN.

Bewley, T. F. (1983). A difficulty with the optimum quantity of money. *Econometrica*, 51(5):1485–1504.

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2017a). Fiscal policy and debt management with incomplete markets. *The Quarterly Journal of Economics*, 132(2):617–663.

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2017b). Public debt in economies with heterogeneous agents. *Journal of Monetary Economics*, 91:39–51.

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

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2026). Efficiency, insurance, and redistribution effects of government policies. Working paper, February 2026.

Bhattacharya, J., Haslag, J. H., and Martin, A. (2005). Heterogeneity, redistribution, and the Friedman rule. *International Economic Review*, 46(2):437–454.

Bianchi, J. and Bigio, S. (2022). Banks, liquidity management, and monetary policy. *Econometrica*, 90(1):391–454.

Bigio, S. and Sannikov, Y. (2026). A model of credit, money, interest, and prices. *The Journal of Finance*. Forthcoming; NBER Working Paper 28540.

Board of Governors of the Federal Reserve System (2024). Federal reserve banks combined financial statements as of and for the years ended december 31, 2023 and 2022. Technical report, Board of Governors of the Federal Reserve System.

Boel, P. and Camera, G. (2006). Efficient monetary allocations and the illiquidity of bonds. *Journal of Monetary Economics*, 53:1693–1715.

Bohn, H. (1998). The behavior of U.S. public debt and deficits. *Quarterly Journal of Economics*, 113(3):949–963.

Bustamante, C. (2023). The long-run redistributive effects of monetary policy. *Journal of Monetary Economics*, 140:106–123.

Calvo, G. A. (1978). On the time consistency of optimal policy in a monetary economy. *Econometrica*, 46(6):1411–1428.

Challe, E. (2020). Uninsured unemployment risk and optimal monetary policy in a zero-liquidity economy. *American Economic Journal: Macroeconomics*, 12(2):241–283.

Chamley, C. and Polemarchakis, H. M. (1984). Assets, general equilibrium and the neutrality of money. *Review of Economic Studies*, 51(1):129–138.

Chang, R. (1998). Credible monetary policy in an infinite horizon model: Recursive approaches. *Journal of Economic Theory*, 81(2):431–461.

Chari, V. V., Christiano, L. J., and Kehoe, P. J. (1996). Optimality of the Friedman rule in economies with distorting taxes. *Journal of Monetary Economics*, 37(2):203–223.

Chari, V. V. and Kehoe, P. J. (1999). Optimal fiscal and monetary policy. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1C, chapter 26, pages 1671–1745. Elsevier.

Chien, Y. and Wen, Y. (2024). The Ramsey steady-state conundrum in heterogeneous-agent economies. *Journal of Economic Theory*, 220:105873.

Copeland, A., Duffie, D., and Yang, Y. D. (2025). Reserves were not so ample after all. *Quarterly Journal of Economics*, 140(1):239–281.

Cúrdia, V. and Woodford, M. (2011). The central-bank balance sheet as an instrument of monetary policy. *Journal of Monetary Economics*, 58(1):54–79.

da Costa, C. E. and Werning, I. (2008). On the optimality of the Friedman rule with heterogeneous agents and nonlinear income taxation. *Journal of Political Economy*, 116(1):82–112.

Dávila, E. and Schaab, A. (2023). Optimal monetary policy with heterogeneous agents: Discretion, commitment, and timeless policy. NBER Working Paper 30961, February 2023.

Del Negro, M. and Sims, C. A. (2015). When does a central bank's balance sheet require fiscal support? *Journal of Monetary Economics*, 73:1–19.

D'Erasmo, P., Mendoza, E. G., and Zhang, J. (2016). What is a sustainable public debt? In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 2493–2597. Elsevier.

Díaz-Giménez, J., Giovannetti, G., Marimon, R., and Teles, P. (2008). Nominal debt as a burden on monetary policy. *Review of Economic Dynamics*, 11(3):493–514.

Duffie, D. (2026). The payment system puts a floor on the fed's balance sheet. *Brookings Papers on Economic Activity*. Spring, conference draft.

Friedman, M. (1969). The optimum quantity of money. In *The Optimum Quantity of Money and Other Essays*, pages 1–50. Aldine Publishing Company, Chicago.

Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34.

Gorton, G. and Ordoñez, G. (2022). The supply and demand for safe assets. *Journal of Monetary Economics*, 125:132–147.

Greenwood, R., Hanson, S. G., and Stein, J. C. (2015). A comparative-advantage approach to government debt maturity. *The Journal of Finance*, 70(4):1683–1722.

Hagedorn, M. (2017). A demand theory of the price level. Working paper, University of Oslo and CEPR, August 2017.

Hagedorn, M. (2024a). A comment on "Paper Money" by Chris Sims. Working paper, University of Oslo and CEPR, November 2024.

Hagedorn, M. (2024b). The failed theory of the price level. Working paper, University of Oslo and CEPR, January 2024.

Hall, R. E. and Reis, R. (2015). Maintaining central-bank financial stability under new-style central banking. Working Paper 21173, National Bureau of Economic Research.

Herrenbrueck, L. (2019). Frictional asset markets and the liquidity channel of monetary policy. *Journal of Economic Theory*, 181:82–120.

Holmström, B. and Tirole, J. (1998). Private and public supply of liquidity. *Journal of Political Economy*, 106(1):1–40.

Hurtado, S., Nuño, G., and Thomas, C. (2023). Monetary policy and sovereign debt sustainability. *Journal of the European Economic Association*, 21(1):293–325.

Ireland, P. N. (2009). On the welfare cost of inflation and the recent behavior of money demand. *American Economic Review*, 99(3):1040–1052.

Kaplan, G. (2025). Implications of fiscal-monetary interaction from HANK models. Working Paper 34117, National Bureau of Economic Research.

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

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Kehoe, T. J., Levine, D. K., and Woodford, M. (1990). The optimum quantity of money revisited. Working Paper 404, Federal Reserve Bank of Minneapolis.

Kimbrough, K. P. (1986). The optimum quantity of money rule in the theory of public finance. *Journal of Monetary Economics*, 18(3):277–284.

Klein, P., Krusell, P., and Ríos-Rull, J.-V. (2008). Time-consistent public policy. *The Review of Economic Studies*, 75(3):789–808.

Kocherlakota, N. R. (2003). Societal benefits of illiquid bonds. *Journal of Economic Theory*, 108:179–193.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2012). The aggregate demand for Treasury debt. *Journal of Political Economy*, 120(2):233–267.

Kydland, F. E. and Prescott, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491.

Lagos, R. and Navarro, G. (2023). Monetary policy operations: Theory, evidence, and tools for quantitative analysis. NBER Working Paper 31370.

Lagos, R., Rocheteau, G., and Wright, R. (2017). Liquidity: A New Monetarist perspective. *Journal of Economic Literature*, 55(2):371–440.

Lagos, R. and Wright, R. (2005). A unified framework for monetary theory and policy analysis. *Journal of Political Economy*, 113(3):463–484.

Le Grand, F., Martin-Baillon, A., and Ragot, X. (2022). Should monetary policy care about redistribution? Optimal monetary and fiscal policy with heterogeneous agents. *The Review of Economic Studies*. Forthcoming; December 2022 accepted manuscript.

Leeper, E. M. (1991). Equilibria under 'active' and 'passive' monetary and fiscal policies. *Journal of Monetary Economics*, 27(1):129–147.

Lippi, F., Ragni, S., and Trachter, N. (2015). Optimal monetary policy with heterogeneous money holdings. *Journal of Economic Theory*, 159:339–368. Part A.

Lucas, Robert E., J. and Nicolini, J. P. (2015). On the stability of money demand. *Journal of Monetary Economics*, 73:48–65.

Lucas, R. E. and Stokey, N. L. (1983). Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics*, 12(1):55–93.

Marcet, A. and Marimon, R. (2019). Recursive contracts. *Econometrica*, 87(5):1589–1631.

Martin, A. and Monnet, C. (2011). Monetary policy implementation frameworks: A comparative analysis. *Macroeconomic Dynamics*, 15(S1):145–189.

Martin, F. M. (2009). A positive theory of government debt. *Review of Economic Dynamics*, 12(4):608–631.

Martin, F. M. (2011). On the joint determination of fiscal and monetary policy. *Journal of Monetary Economics*, 58:132–145.

Martin, F. M. (2013). Government policy in monetary economies. *International Economic Review*. Also Federal Reserve Bank of St. Louis Working Paper 2011-026B.

Martin, F. M. (2015). Debt, inflation and central bank independence. *European Economic Review*, 79:129–150.

Molico, M. (2006). The distribution of money and prices in search equilibrium. *International Economic Review*, 47(3):701–722.

Nagel, S. (2016). The liquidity premium of near-money assets. *Quarterly Journal of Economics*, 131(4):1927–1971.

Niemann, S., Pichler, P., and Sorger, G. (2013). Public debt, discretionary policy, and inflation persistence. *Journal of Economic Dynamics and Control*, 37(6):1097–1109.

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

Obstfeld, M. (1997). Dynamic seigniorage theory: An exploration. *Macroeconomic Dynamics*, 1(3):588–614.

Obstfeld, M. and Rogoff, K. (1983). Speculative hyperinflations in maximizing models: Can we rule them out? *Journal of Political Economy*, 91(4):675–687.

Persson, M., Persson, T., and Svensson, L. E. O. (2006). Time consistency of fiscal and monetary policy: A solution. *Econometrica*, 74(1):193–212.

Phelps, E. S. (1973). Inflation in the theory of public finance. *Swedish Journal of Economics*, 75(1):67–82.

Roberds, W. (1987). Models of policy under stochastic replanning. *International Economic Review*, 28(3):731–755.

Rocheteau, G., Weill, P.-O., and Wong, T.-N. (2018a). A tractable model of monetary exchange with ex post heterogeneity. *Theoretical Economics*, 13(3):1155–1195. Also NBER WP 21179 (2015), Federal Reserve Bank of Richmond WP 17-06 (2017).

Rocheteau, G. and Wright, R. (2005). Money in search equilibrium, in competitive equilibrium, and in competitive search equilibrium. *Econometrica*, 73(1):175–202.

Rocheteau, G., Wright, R., and Xiao, S. X. (2018b). Open market operations. *Journal of Monetary Economics*, 98:114–128.

Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *The Quarterly Journal of Economics*, 100(4):1169–1189.

Sargent, T. J. (1982). The ends of four big inflations. In Hall, R. E., editor, *Inflation: Causes and Effects*, pages 41–97. University of Chicago Press, Chicago.

Sargent, T. J. and Wallace, N. (1975). 'rational' expectations, the optimal monetary instrument, and the optimal money supply rule. *Journal of Political Economy*, 83(2):241–254.

Sargent, T. J. and Wallace, N. (1981). Some unpleasant monetarist arithmetic. *Federal Reserve Bank of Minneapolis Quarterly Review*, 5(3):1–17.

Schaumburg, E. and Tambalotti, A. (2007). An investigation of the gains from commitment in monetary policy. *Journal of Monetary Economics*, 54(2):302–324.

Schulhofer-Wohl, S. (2025). Monetary policy implementation and the consolidated government balance sheet. Federal Reserve Bank of Dallas. April 15, 2025.

Shi, S. (2008). Efficiency improvement from restricting the liquidity of nominal bonds. *Journal of Monetary Economics*, 55:1025–1037.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15(2):127–148.

Sims, C. A. (1994). A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory*, 4(3):381–399.

Sterk, V. and Tenreyro, S. (2018). The transmission of monetary policy through redistributions and durable purchases. *Journal of Monetary Economics*, 99:124–137.

Straub, L. and Werning, I. (2020). Positive long-run capital taxation: Chamley–Judd revisited. *American Economic Review*, 110(1):86–119.

U.S. Bureau of Economic Analysis (2024). Gross domestic product, fourth quarter and year 2023 (second estimate). Technical report, U.S. Department of Commerce.

Wallace, N. (1981). A Modigliani-Miller theorem for open-market operations. *American Economic Review*, 71(3):267–274.

Walsh, C. E. (1995). Optimal contracts for central bankers. *American Economic Review*, 85(1):150–167.

Williamson, S. D. (2025). Interest rate control and interbank markets. Working paper, University of Western Ontario, July 2025.

Woodford, M. (1990). Public debt as private liquidity. *American Economic Review*, 80(2):382–388. Papers and Proceedings.

Woodford, M. (1995). Price-level determinacy without control of a monetary aggregate. *Carnegie-Rochester Conference Series on Public Policy*, 43:1–46.

A Proofs of Propositions and Lemmas

Throughout these proofs we use the *adjoint identity*: for the household generator $\mathcal{A}_s$ and its adjoint $\mathcal{A}_s^*$ (Appendix B) and any smooth γ_s,

$$\int \gamma_s(x)\, \mathcal{A}_s^*[g_s]\, dx = \int \mathcal{A}_s[\gamma_s](x) \cdot g_s(x)\, dx, \tag{66}$$

where $\mathcal{A}_s[\gamma_s] = s_s^\omega\, \partial_\omega \gamma_s + \alpha[\gamma_s(\omega - y_s, z) - \gamma_s] + \mathcal{A}_z[\gamma_s]$ and $s_s^\omega = r_s^a \omega + i_s^b(\omega - y_s) + z_s - c_s + \mathcal{T}_s$ is the wealth savings rate (using $a_s = y_s$, $b_s = \omega - y_s$). It moves the adjoint generator off the density onto the social value γ_s by integration by parts.

A.1 Proof of Lemma 1

We establish the viscosity solution property, boundary condition, monotonicity, and concavity results on the state space $[0, \infty) \times \mathcal{Z}$. Define the HJB operator $\mathcal{H}_t : [0, \infty) \times \mathcal{Z} \times \mathbb{R} \to \mathbb{R}$ by

$$\begin{aligned} \mathcal{H}_t(\omega, z, p_\omega) \;\equiv\; & \rho W_t - u(c_t) - \alpha U(y_t) \\ & - [r_t^b \omega - i_t^b y_t + z_t - c_t + \mathcal{T}_t]\, p_\omega \\ & - \alpha[W_t(\omega - y_t, z) - W_t(\omega, z)] - \mathcal{A}_z W_t - \partial_t W_t, \end{aligned} \tag{67}$$

where we have already substituted the household's optimal $a_t = y_t$ (Section 2.3), so the HJB equation (14) reads $\mathcal{H}_t(\omega, z, \partial_\omega W_t) = 0$.

A continuous function v is a *constrained viscosity subsolution* of (14) if, for each $(\omega_0, z_0) \in (0, \infty) \times \mathcal{Z}$ and each smooth ϕ such that $v - \phi$ attains a local maximum at (ω_0, z_0) with $v(\omega_0, z_0) = \phi(\omega_0, z_0)$:

$$\mathcal{H}_t(\omega_0, z_0, \partial_\omega \phi(\omega_0)) \leq 0. \tag{68}$$

A continuous function v is a *constrained viscosity supersolution* of (14) if, for each $(\omega_0, z_0) \in [0, \infty) \times \mathcal{Z}$ (including the boundary $\omega_0 = \underline{b}$) and each smooth ϕ such that $v - \phi$ attains a local minimum at (ω_0, z_0) with $v(\omega_0, z_0) = \phi(\omega_0, z_0)$:

$$\mathcal{H}_t(\omega_0, z_0, \partial_\omega \phi(\omega_0)) \geq 0. \tag{69}$$

Note the asymmetry: the subsolution condition is required only on the interior, while the supersolution condition is required on the full domain including the boundary $\omega_0 = \underline{b}$. This is the standard definition for state-constrained viscosity solutions; see Achdou et al. (2022, Appendix D).

Part (i): W_t is a constrained viscosity solution. Fix $(\omega_0, z_0) \in (0, \infty) \times \mathcal{Z}$ and let ϕ be a smooth

test function such that $W_t - \phi$ attains a local maximum at (ω_0, z_0) with $W_t(\omega_0, z_0) = \phi(\omega_0, z_0)$. By the Dynamic Programming Principle applied to W_t,[40] for any $h > 0$,

$$W_t(\omega_0, z_0) = \mathbb{E}_{\omega_0, z_0}\left[\int_t^{t+h} e^{-\rho(s-t)}\mathcal{U}_s^*\, ds + e^{-\rho h} W_{t+h}(\omega_{t+h}^*, z_{t+h})\right].$$

Applying Dynkin's formula to $e^{-\rho s}\phi(\omega_s, z_s)$ over $[t, t+h]$, using the local-maximum property of ϕ on continuous and Poisson-jump segments, dividing by h, and sending $h \to 0$ yields the subsolution condition (68). The supersolution argument is symmetric, using the DPP inequality $W_t(\omega_0, z_0) \geq \mathbb{E}[\cdots]$ for any admissible control and applying Dynkin's formula to ϕ under $W_t \geq \phi$ throughout. The argument extends to the boundary $\omega_0 = \underline{b}$ by restricting the admissible-control set to policies that maintain $\omega_s \geq \underline{b}$.

Part (ii): Boundary condition at $\omega = 0$. At $\omega = 0$, the constraints $a_t \geq 0$ and $b_t \geq \underline{b}$ with $a_t + b_t = \omega = 0$ force $a_t = 0$ and $b_t = \underline{b}$, and the constraint $y_t \leq a_t$ then forces $y_t(\underline{b}, z) = 0$. Assumption 1 guarantees positive resources net of debt-service costs: $\underline{z}_t + \mathcal{T}_t(\underline{b}, \underline{z}) + r_t^b \underline{b} > 0$ for any $r_t^b \in [0, \rho]$, hence $z + \mathcal{T}_t(\underline{b}, z) + r_t^b \underline{b} > 0$ for all z. The Inada condition $u'(0) = \infty$ ensures $c_t(\underline{b}, z) > 0$. The Euler condition gives $\partial_\omega W_t(\underline{b}, z) = u'(c_t(\underline{b}, z)) < u'(0) = \infty$, so the marginal value at the boundary is finite. The state constraint $\dot{\omega}_t \geq 0$ at $\omega = 0$, using $a_t = y_t = 0$, gives $\dot{\omega}_t = r_t^b \underline{b} + z + \mathcal{T}_t(\underline{b}, z) - c_t \geq 0$, i.e., $c_t \leq z + \mathcal{T}_t(\underline{b}, z) + r_t^b \underline{b}$. Since $u'' < 0$,

$$\partial_\omega W_t(\underline{b}, z) = u'(c_t) \geq u'(z + \mathcal{T}_t(\underline{b}, z) + r_t^b \underline{b}),$$

establishing (19). Equality holds iff $\dot{\omega}_t = 0$.

Part (iii): Monotonicity and concavity. Fix t, z, $\lambda \in (0,1)$, and $\omega^0 \neq \omega^1$ with $\hat{\omega} = \lambda\omega^0 + (1-\lambda)\omega^1$. Let $\pi^i = (c_s^i, y_s^i)_{s \geq t}$ denote the optimal plan from (ω^i, z), and define the mixture plan $\hat{\pi} \equiv \lambda\pi^0 + (1-\lambda)\pi^1$. Since the wealth budget (4) is affine in (ω, c, y) for given z, the mixture wealth $\hat{\omega}_s \equiv \lambda\omega_s^0 + (1-\lambda)\omega_s^1$ satisfies the budget constraint under $\hat{\pi}$. By strict concavity of u and U, the mixture plan yields strictly higher utility than the convex combination of π^0 and π^1 on a set of positive measure (since $\omega^0 \neq \omega^1$ implies the paths differ near t). Therefore $W_t(\hat{\omega}, z) > \lambda W_t(\omega^0, z) + (1-\lambda) W_t(\omega^1, z)$. Strict monotonicity follows from the strict relaxation of the budget constraint by additional wealth: more ω allows strictly more consumption (flow or lumpy) or strictly more saving, both of which raise W_t.

Part (iv): Asymptotic $\partial_\omega W_t \to 0$. As $\omega \to \infty$, the marginal lumpy benefit $\alpha[U'(y) - W_\omega(\omega - y, z)] \to 0$ (since $U'(y)$ stays bounded and $W_\omega(\omega - y, z) \to 0$ by the same logic). The lumpy FOC (13) then implies $i_t^b \partial_\omega W_t(\omega, z) \to 0$, which gives $\partial_\omega W_t(\omega, z) \to 0$ for $i_t^b > 0$. (For $i_t^b = 0$, the

[40] Invoking the DPP and Dynkin's formula requires that the optimal drift $|\dot{\omega}_s^*|$ is bounded along optimal paths. The drift $r_s^b \omega - i_s^b y^* + z_s - c^* + \mathcal{T}_s$ is bounded because c^* and y^* are finite, extending the regularity results of Rocheteau et al. (2018a, Lemmas 1–2) to the present setting.

asymptotic limit follows from the Inada condition $u'(\infty) = 0$ applied to the consumption Euler.) □

A.2 Proof of Lemma 2

The household chooses $y \in [0, \omega]$ to maximize the right-hand side of the HJB at state (ω, z). The upper bound is the maximum money the household can hold given the debt limit $b \geq \underline{b}$: with $a = y$ (binding when $i^b > 0$) and $b = \omega - a \geq \underline{b}$, we get $y \leq \omega$. Substituting the optimal $a = y$ (which holds whenever $i^b > 0$; see Section 2.3), the lumpy choice problem reduces to

$$\max_{0 \leq y \leq \omega} \{\alpha[U(y) + W_t(\omega - y, z) - W_t(\omega, z)] - i_t^b y \cdot \partial_\omega W_t(\omega, z)\}.$$

The objective is strictly concave in y by strict concavity of U and of $W_t(\cdot, z)$ (Lemma 1(iii)), so the maximizer is unique. The Karush-Kuhn-Tucker conditions give the lumpy FOC (13): $\alpha[U'(y) - \partial_\omega W_t(\omega - y, z)] \geq i_t^b \partial_\omega W_t(\omega, z)$, with equality whenever $y < \omega$ and inequality (slack) at the corner $y = \omega$.

Part (i): Existence of the full-depletion threshold. The corner $y = \omega$ is optimal iff

$$\alpha[U'(\omega) - \partial_\omega W_t(\underline{b}, z)] \geq i_t^b \partial_\omega W_t(\omega, z),$$

i.e., the marginal lumpy benefit at the full-depletion endpoint exceeds the marginal cost of holding wealth in money form. (After full debt-financed depletion at ω, the household's post-shock state is $\omega - y = \underline{b}$.) By the Inada condition $U'(0) = \infty$, $\partial_\omega W_t(\underline{b}, z) < \infty$ (Lemma 1(ii)), and the asymptotic $\partial_\omega W_t \to 0$ as $\omega \to \infty$ (Lemma 1(iv)), the function

$$\Delta_t(\omega; z) \equiv \alpha[U'(\omega) - \partial_\omega W_t(\underline{b}, z)] - i_t^b \partial_\omega W_t(\omega, z)$$

satisfies $\Delta_t(\underline{b}^+; z) = +\infty$ (since $U'(\omega) \to \infty$ as $\omega \to \underline{b}^+$) and $\Delta_t(\infty; z) < 0$ (negative because $U'(\infty) - \partial_\omega W_t(\underline{b}, z) < 0$ when $\partial_\omega W_t(\underline{b}, z) > 0$). By the intermediate value theorem and the strict monotonicity of Δ_t in ω (which follows from $U''(\cdot) < 0$ and $\partial_{\omega\omega} W_t < 0$), there exists a unique $\underline{\omega}_t(z) > 0$ solving $\Delta_t(\underline{\omega}_t(z); z) = 0$, which is precisely (20). For $\omega \in (\underline{b}, \underline{\omega}_t(z)]$, $\Delta_t(\omega; z) \geq 0$, so the corner $y = \omega$ is optimal: $y_t(\omega, z) = \omega$. For $\omega > \underline{\omega}_t(z)$, $\Delta_t(\omega; z) < 0$, the corner is not optimal, and the interior FOC (13) holds with equality, giving $y_t(\omega, z) < \omega$.

Part (ii): Monotonicity $0 \leq \partial_\omega y_t \leq 1$**.** For $\omega \in (\underline{b}, \underline{\omega}_t(z)]$, $y_t(\omega, z) = \omega$, so $\partial_\omega y_t = 1$. For $\omega > \underline{\omega}_t(z)$, the interior FOC $\alpha[U'(y_t) - \partial_\omega W_t(\omega - y_t, z)] = i_t^b \partial_\omega W_t(\omega, z)$ holds. Implicit differentiation with respect to ω:

$$\alpha[U''(y_t) + \partial_{\omega\omega} W_t(\omega - y_t, z)]\partial_\omega y_t - \alpha \partial_{\omega\omega} W_t(\omega - y_t, z) = i_t^b \partial_{\omega\omega} W_t(\omega, z).$$

Solving:

$$\partial_\omega y_t = \frac{\alpha \partial_{\omega\omega} W_t(\omega - y_t, z) + i_t^b \partial_{\omega\omega} W_t(\omega, z)}{\alpha[U''(y_t) + \partial_{\omega\omega} W_t(\omega - y_t, z)]}.$$

Both numerator and denominator are negative (by concavity of U and W_t), so $\partial_\omega y_t > 0$. Moreover, $|U''(y_t)| > 0$ implies the denominator's absolute value exceeds the absolute value of the numerator's first term; the second term in the numerator $i_t^b \partial_{\omega\omega} W_t(\omega, z)$ adds a small magnitude. So $\partial_\omega y_t < 1$. □

A.3 Proof of Lemma 3

Fix $g^M = \pi > -\rho$, so $r^a = -g^M$ and $r^b = r^a + i^b = i^b - \pi$. What is left is to argue there exists a liquidity premium such that both markets clear at the given inflation rate. Prices and the transfer enter the household problem as constants. Steps 1 through 3 treat both as parameters, and the budget determines the transfer in Step 4. For i^b in $\mathcal{I} \equiv (0, \rho + g^M)$ the state $x = (\omega, z)$ follows a piecewise-deterministic Markov process: between jumps wealth drifts at $\dot{\omega}(\omega, z) = r^b \omega - i^b y(\omega, z) + z - c(\omega, z) + T$, income switches between states at rates λ_f and λ_s, and at rate α wealth jumps from ω to $\omega - y(\omega, z)$.

Step 1: policies are jointly continuous in i^b. Let $i_n^b \to i_\infty^b$ in $\mathcal{I}$. The coefficients of the HJB (14) converge locally uniformly — uniformly on every compact subset of $[0, \infty) \times \mathcal{Z}$ — so by stability of constrained viscosity solutions together with a comparison principle for state-constrained problems of this class (see Achdou et al., 2022), $W_n \to W_\infty$ locally uniformly on $[0, \infty) \times \mathcal{Z}$. Each $W_n(\cdot, z)$ is strictly concave (Lemma 1(iii)) and the limit is differentiable on $(0, \infty)$, so convergence carries to the derivatives $\partial_\omega W_n \to \partial_\omega W_\infty$ locally uniformly. The consumption Euler equation and strict monotonicity of u' give $c_n \to c_\infty$ locally uniformly. On the interior region the left side of the lumpy condition (13) is strictly decreasing in y and its right side is continuous, so y is the unique root and converges; on the depletion region $y = \omega$; and the threshold $\underline{\omega}(z)$, the unique crossing in (20), converges as well. Policies are therefore continuous in (ω, z, i^b) on $[0, \infty) \times \mathcal{Z} \times \mathcal{I}$. The same argument applies to the constant transfer, so continuity holds jointly in the premium and the transfer.

Step 2: each premium admits a unique stationary distribution. First, a drift bound at the top. Since $y \geq 0$, the mean drift of wealth satisfies $\dot{\omega}(\omega, z) - \alpha y(\omega, z) \leq r^b \omega + \bar{z} + T - c(\omega, z)$. With $r^b < \rho$, consumption is asymptotically bounded below relative to wealth, $\liminf_{\omega \to \infty} c(\omega, z)/\omega \geq r^b + (\rho - r^b)/\sigma$. The argument adapts the asymptotic analysis of the one-asset CRRA problem in Achdou et al. (2022): differentiating the HJB and using both first-order conditions, the marginal value $m = \partial_\omega W$ satisfies a linear equation whose lumpy terms preserve the one-asset decay bound, and the consumption asymptotics follow from $c = (u')^{-1}(m)$. Hence there exists ω^* such that the flow drift $\dot{\omega}(\omega, z)$ is strictly negative for all $\omega \geq \omega^*$ in both income states. Because lumpy shocks only reduce wealth, every path enters the interval $[0, \omega^*]$ in finite time and never leaves

it. Any stationary distribution therefore has support in $[0, \omega^*] \times \mathcal{Z}$, and existence follows from the standard argument on a compact state space.

Step 3: the stationary distribution and its aggregates are continuous in the premium and the transfer. By Step 1 the confinement bound ω^* can be chosen uniformly on compact sets of premia and transfers, so along $(i^b_n, T_n) \to (i^b_\infty, T_\infty)$ the invariant measures have uniformly bounded supports. Let G' be any weak limit point. The extended generators converge on smooth test functions by Step 1, so G' is invariant for the limit process, and uniqueness in Step 2 forces G' to be the invariant measure at (i^b_∞, T_∞). With uniformly bounded supports, weak convergence directly gives convergence of the aggregates, so money demand $Q(i^b, T)$, wealth $A(i^b, T)$, and the demanded ratio $D(i^b, T) = (A - Q)/Q$ are continuous in the pair, with $Q > 0$ from $U'(0) = \infty$.

Step 4: a budget-consistent transfer at each premium. Integrating the wealth law of motion against the stationary distribution, stationarity of mean wealth, jump terms included, gives the accounting identity

$$\int c \, dG \;+\; \alpha \int y \, dG \;=\; \bar{Y} \;+\; T \;+\; r^b A(i^b, T) \;-\; i^b Q(i^b, T), \tag{70}$$

where $\bar{Y} \equiv \int z \, dG$ is aggregate income. Using $r^b = i^b - \pi$ and $A^b = A - Q$, the budget (7) holds at the demanded aggregates exactly when

$$T \;=\; i^b \, Q(i^b, T) \;-\; r^b A(i^b, T). \tag{71}$$

Define $H(T) \equiv T - i^b Q(i^b, T) + r^b A(i^b, T)$, continuous on the feasible interval by Step 3, and let $\mathcal{I}^* \subseteq \mathcal{I}$ collect the premia at which H changes sign, so that a budget-consistent transfer $T(i^b)$ exists by the intermediate value theorem. The first term of H has unit slope in the transfer, while the remaining terms are scaled by the rates i^b and r^b, both below $\rho + g^M$. A change of sign can therefore fail only if the demand for money or wealth responds to the transfer by more than the reciprocal of those rates. When the consistent transfer is unique at each premium in $\mathcal{I}^*$, its graph is closed by Step 3 and lies in a compact set, so $T(i^b)$ is continuous and the relative demand along the equilibrium branch, $D(i^b) \equiv D(i^b, T(i^b))$, is continuous. If several transfers are consistent at some premium, the construction selects one, and the same discussion gives the condition for uniqueness.

Step 5: market clearing. As $i^b \downarrow 0$ the policies converge to their zero-premium limits, where $a = y$ is still imposed, and Steps 2–4 apply at $i^b = 0$, so $D(0^+) = D(0) \equiv \underline{\beta}(g^M) \in (0, \infty)$. Define $\bar{\beta}(g^M) \equiv \max\{\sup_{i^b \in \mathcal{I}^*_0} D(i^b), \underline{\beta}(g^M)\}$. For any $\beta \in (\underline{\beta}(g^M), \bar{\beta}(g^M))$ there are premia at which $D < \beta$ (near zero) and premia at which $D > \beta$, so by continuity of D and the intermediate value theorem the bond market clears at some $i^b(\beta) \in \mathcal{I}^*_0$, and the premium is strictly positive. For $\beta \leq \underline{\beta}(g^M)$, set $i^b = 0$: money and bonds are perfect substitutes, each household is indifferent

over money holdings in $[y, \omega]$, and aggregate money demand can be assigned any value in $[Q(0), A(0)]$; since $\beta \leq \underline{\beta}(g^M)$ is equivalent to $A(0)/(1+\beta) \geq Q(0)$, the composition is absorbed along the indifference with the same real allocation. □

A.4 Proof of Proposition 1

Proposition 1 is proved in Appendix C, which derives the full optimality system—both the commitment ($\lambda \to 0$) and discretion ($\lambda \to \infty$) limits—by the discretize-first method of Dávila and Schaab (2023) in detrended-nominal variables, with the joint continuum limit recovering the conditions of Proposition 1.

A.5 Proof of Lemma 4

The proof has two steps. Step 1 establishes that $\psi_s = 0$ is *self-sustaining*: whenever it holds, all source terms in the forward KFE (33) vanish, so $\psi_s = 0$ remains consistent going forward. Step 2 establishes the initial condition $\psi_0 = 0$ under each regime, and then invokes forward uniqueness of the linear PDE to conclude $\psi_s \equiv 0$ for all s.

Step 1 (Self-sustaining property). Suppose $\psi_s(x) = 0$ at date s. The transfer FOC (27) immediately gives $\partial_\omega \gamma_s(x) = \theta_s$ for all x: the social shadow value of wealth is uniform across the state space. We now show this forces $\chi_s = \eta_s = 0$.

Flow consumption (χ_s): From (31), the flow consumption wedge is $\tilde{\chi}_s \equiv \chi_s u''(c_s)/g_s = \partial_\omega \gamma_s - u'(c_s)$. Here $\partial_\omega \gamma_s = \theta_s$ is the planner's fiscal shadow price—the social value of a marginal unit of wealth at x—while $u'(c_s)$ is the household's marginal utility of consumption at that point. With an unrestricted transfer, the planner can independently adjust the consumption of each agent x by varying $\mathcal{T}_s(x)$. Optimality of the transfer therefore requires the planner to equate these: $u'(c_s) = \partial_\omega \gamma_s = \theta_s$, so $\chi_s = 0$.

Lumpy consumption (η_s): From (32), the lumpy consumption wedge is $\tilde{\eta}_s = \alpha[\partial_\omega \gamma_s(\omega - y_s, z) - U'(y_s)] + i_s^b[\partial_\omega \gamma_s(\omega, z) - \theta_s]$. With $\partial_\omega \gamma_s = \theta_s$ uniform, the second (portfolio) term vanishes. For the first (lumpy) term: with an unrestricted transfer, the planner uses $\mathcal{T}_s$ to implement $U'(y_s) = \partial_\omega \gamma_s(\omega - y_s, z) = \theta_s$ for all x, so $\tilde{\eta}_s = 0$ and $\eta_s = 0$.

With $\chi_s = \eta_s = 0$ at every x, all source terms in (33) vanish. The KFE reduces to $\partial_t \psi_s = \mathcal{A}_s^*[\psi_s]$, which is satisfied by $\psi_s = 0$.

Step 2 (Initial conditions and uniqueness). Step 1 shows that $\psi_s = 0$ is self-sustaining. It remains to establish $\psi_0 = 0$ in each case.

Full discretion ($\lambda \to \infty$). $\psi_s \equiv 0$ by definition.

Full commitment ($\lambda \to 0$), standard Ramsey. The initial condition in Proposition 1 gives $\psi_0 = 0$. With zero source terms (Step 1), the forward KFE propagates $\psi_s = 0$ uniquely.

Full commitment ($\lambda \to 0$), stationary equilibrium. The stationary ψ_{ss} satisfies $0 = \mathcal{A}_{ss}^*[\psi_{ss}] +$ (source terms). Step 1 shows the source terms vanish when $\psi_{ss} = 0$, so $\mathcal{A}_{ss}^*[\psi_{ss}] = 0$, whose unique solution is $\psi_{ss} = 0$.

In all cases $\psi_s \equiv 0$. Substituting into (27) yields (38). □

A.6 Proof of Proposition 2

By Lemma 4, $\psi_s \equiv 0$ under both full discretion and full commitment with unrestricted transfers. The optimality conditions therefore coincide across regimes.

With $\psi_s = 0$, the transfer FOC (38) gives $\partial_\omega \gamma_s = \theta_s$ for all x. The planner uses $\mathcal{T}_s(x)$ to enforce $u'(c_s(x)) = \theta_s$ for all x, so consumption is equalized across all wealth levels and $\partial_\omega c_s = 0$. Differentiating the household HJB with respect to ω and applying the envelope theorem gives $\dot{\theta}_s = (\rho - r_s^a)\theta_s$. Combined with the money supply FOC (29), $\dot{\theta}_s = (\rho - r_s^a)\theta_s - \mu_s$, this requires $\mu_s = 0$.

From (36): $\mu_s - \varphi_s = i_s^b \theta_s$, so $\varphi_s = -i_s^b \theta_s$. The bond supply FOC (30), $\dot{\theta}_s = (\rho - r_s^b)\theta_s - \varphi_s$, combined with $r_s^b = r_s^a + i_s^b$ and $\mu_s = 0$, is automatically consistent. With consumption and lumpy allocations equalized across x, the planner has no remaining motive to drive a wedge between the returns on money and bonds, so $i_s^b = 0$ and hence $\varphi_s = 0$.

Since $u'(c_s) = \theta_s$ holds at all x, the equilibrium allocations are determined pointwise by θ_s alone: $c_s = c^*(\theta_s)$ and $y_s = y^*(\theta_s)$. Goods-market clearing requires $c^*(\theta_s) + \alpha\, y^*(\theta_s) = \mathbb{E}_z[z]$, which uniquely determines $\theta_s = \theta^{FB}$ (time invariant). Since $\dot{\theta}_s = 0$, $(\rho - r_s^a)\theta^{FB} = 0$ forces $r_s^a = \rho$. The conditions (2)–(3) are satisfied.

Finally, $B_t = 0$ is optimal: with $r_t^a = r_t^b = \rho$ ($i_t^b = 0$), money and bonds are perfect substitutes in return; bonds lack only the lumpy-financing role of money. Each household is indifferent across any portfolio composition with at least $y^*(\theta^{FB})$ held in money, so an equilibrium with $B_t = 0$—in which all wealth is held as money—is consistent with optimal household behavior and clears both markets when $\phi_t M_t = \int y^*(\theta^{FB})\, dG_t$. □

A.7 Proof of Proposition 3

Under full discretion ($\psi_s \equiv 0$) in steady state ($\partial_t \gamma = \partial_t W = \partial_t G = \dot{\theta} = 0$), the optimality system of Proposition 1 reduces to: $\theta = \int \partial_\omega \gamma \, dG$, $\mu = (\rho - r^a)\theta$, $\varphi = (\rho - r^a - i^b)\theta$, and the stationary backward HJB, private HJB, and KFE.

Part: $r_{ss}^{a*} < \rho$. The proof follows Achdou et al. (2022) Proposition 4, adapted to our setting with lumpy consumption. Differentiating the private HJB with respect to ω (envelope condition) and

using $\partial_\omega W = u'(c)$:

$$(\rho - r^b)\, u'(c) = u''(c)\, (\partial_\omega c)\, s^\omega + \underbrace{\mathcal{A}_z[u'(c)]}_{\text{income risk}} + \underbrace{\alpha[u'(c(\omega - y, z)) - u'(c)]}_{\text{lumpy precaution}}, \tag{72}$$

where $s^\omega = r^b\omega - i^b y + z - c + \mathcal{T}$ is the wealth drift (with the household's optimal $a = y$ substituted). Set $r^a = \rho$ (Friedman rule for money), so $r^b = \rho + i^b \geq \rho$, and $z = \bar{z}$ (highest productivity). The income-risk and lumpy-precaution terms are strictly positive. With $r^b \geq \rho$ the left-hand side is non-positive, so $u''(c)(\partial_\omega c)s^\omega < 0$; since $u'' < 0$ and $\partial_\omega c > 0$ (Assumption 2), $s^\omega(\omega, \bar{z}; r^a = \rho) > 0$ for all $\omega > \underline{b}$.

As $r^a \uparrow \rho$, the upper bound of the support of the wealth distribution diverges ($\omega_{\max}(r^a) \to \infty$), and aggregate wealth demand $A(r^a) = \int \omega\, dG \to \infty$. A stationary equilibrium requires $A = \phi(M + B)$ finite, so no finite ϕ clears the asset markets at $r^a = \rho$. The Friedman rule for money is therefore infeasible under lump-sum transfers, establishing $r^{a*}_{ss} < \rho$.

Part: Optimal i^{b*}_{ss}. In steady state, the bond supply FOC (30) with $\dot{\theta} = 0$ gives $\varphi_{ss} = (\rho - r^a_{ss} - i^b_{ss})\theta_{ss}$. From the money supply FOC: $\mu_{ss} = (\rho - r^a_{ss})\theta_{ss}$. Subtracting, $\mu_{ss} - \varphi_{ss} = i^b_{ss}\theta_{ss}$. The i^b_{ss} FOC (34) in steady state, with $\psi_{ss} \equiv 0$ under discretion, becomes

$$\int (\omega - y_{ss})\, \partial_\omega \gamma_{ss}\, dG_{ss} \;-\; \int \eta_{ss}\, u'(c_{ss})\, dx \;=\; \theta_{ss}\, \phi_{ss} B_{ss} \;=\; \theta_{ss}\, A^b_{ss},$$

the bond first-order condition in level form. The first term on the left captures the social value of household bond holdings weighted by their portfolio choice $\omega - y_{ss}$; the second is the implementability cost of tightening the lumpy/portfolio FOC by raising i^b; the right-hand side is the fiscal cost of paying interest on outstanding bonds at the shadow price θ_{ss}. Substituting $\theta_{ss} = \int \partial_\omega \gamma_{ss}\, dG_{ss}$ and $A^b_{ss} = \int (\omega - y_{ss})\, dG_{ss}$ (bond-market clearing) and using $\int dG_{ss} = 1$ yields the covariance form (41). □

A.8 Proof of Lemma 5

The reduced social-value equation (42) is linear in γ_{ss}, so its solution decomposes by source. The three sources deliver the level decomposition

$$\gamma_{ss} \;=\; W_{ss} \;+\; \theta_{ss}\, \gamma^f_{ss} \;-\; \frac{\theta_{ss} T_{ss}}{\rho}, \qquad \rho\gamma^f_{ss} \;=\; (\rho - r^a_{ss})\, y_{ss} \;+\; (\rho - r^b_{ss})\, (\omega - y_{ss}) \;+\; \mathcal{A}_{ss}[\gamma^f_{ss}]. \tag{73}$$

First, the household welfare flow $\mathcal{U}_{ss}$: the component $\gamma_{\mathcal{U}}$ solving $\rho\gamma_{\mathcal{U}} = \mathcal{U}_{ss} + \mathcal{A}_{ss}[\gamma_{\mathcal{U}}]$ has the same flow and the same generator as the household's own stationary HJB evaluated at the optimal policy, so $\gamma_{\mathcal{U}} = W_{ss}$, and by the consumption first-order condition $\partial_\omega \gamma_{\mathcal{U}} = \partial_\omega W_{ss} = u'(c_{ss})$.

Second, the terms linear in the household's position: using $\omega = y_{ss} + (\omega - y_{ss})$ and $r^b_{ss} = r^a_{ss} + i^b_{ss}$,

$$(\rho - r^a_{ss})\theta_{ss}\,\omega \;-\; i^b_{ss}\theta_{ss}\,(\omega - y_{ss}) \;=\; \theta_{ss}\Big[(\rho - r^a_{ss})\,y_{ss} + (\rho - r^b_{ss})\,(\omega - y_{ss})\Big],$$

which sources $\theta_{ss}\gamma^f_{ss}$ in (73); the Feynman–Kac representation of γ^f_{ss} in Lemma 5 is immediate from $(\rho - \mathcal{A}_{ss})\gamma^f_{ss} = (\rho - r^a_{ss})y_{ss} + (\rho - r^b_{ss})(\omega - y_{ss})$. Third, the constants: $\rho\theta_{ss}A_{ss} - (\rho - r^a_{ss})\theta_{ss}A_{ss} + i^b_{ss}\theta_{ss}A^b_{ss} = \theta_{ss}[r^a_{ss}A_{ss} + i^b_{ss}A^b_{ss}] = -\theta_{ss}T_{ss}$ by the stationary consolidated budget, and since $\mathcal{A}_{ss}[1] = 0$ a constant source κ contributes κ/ρ to the level of γ_{ss} and nothing to its slope. Summing the three parts gives (73); differentiating in ω gives (43). The reduction uses only the adjoint equation (26)—in which the promise multiplier ψ_{ss} does not appear—the stationary co-state relations $\mu_{ss} = (\rho - r^a_{ss})\theta_{ss}$ and $\varphi_{ss} = (\rho - r^b_{ss})\theta_{ss}$ from (29)–(30) with $\dot\theta_{ss} = 0$, and the stationary budget; the transfer first-order condition, the only optimality condition that differs across regimes, is never invoked. The decomposition therefore holds under discretion and commitment alike, with the regime determining only the value of θ_{ss}. □

A.9 Proof of Proposition 4

We work in the stationary equilibrium of the commitment problem with lump-sum transfers. The key departure from Proposition 3 is that ψ_{ss} is generically nonzero.

Part: Fiscal shadow price. From (28) using $\partial_\omega W_{ss} = u'(c_{ss})$:

$$\theta_{ss} = \int\Big[\partial_\omega\gamma_{ss} + \frac{\psi_{ss}}{g_{ss}}\,u'(c_{ss})\Big]g_{ss}\,dx = \overline{\partial_\omega\gamma_{ss}} + \int \psi_{ss}\,u'(c_{ss})\,dx.$$

Part: Backward HJB and co-state equations. The backward HJB (26) does not contain ψ_{ss} directly. Stationarity and the government budget constraint give $\mu_{ss} = (\rho - r^a_{ss})\theta_{ss}$ and $\varphi_{ss} = (\rho - r^a_{ss} - i^b_{ss})\theta_{ss}$ by the same argument as Proposition 3. These depend only on stationarity and the government budget constraint, not on ψ_{ss}. The backward HJB (42) is therefore structurally identical to the discretion case.

Part: Optimal $r^{a,\text{comm}}_{ss}$ (covariance form). The optimal-inflation condition (35), $R_s = \theta_s A_s$, holds under commitment exactly as under discretion (Appendix C); the regimes differ only through the promise channel $\int \omega\,\psi_{ss}\,\partial_\omega W_{ss}\,dx$ carried in R_{ss}. Evaluated at the stationary commitment equilibrium, $R_{ss} = \theta_{ss}A_{ss}$; subtracting $\theta_{ss}A_{ss}$ and using $\theta_{ss} = \int(\partial_\omega\gamma_{ss} + \frac{\psi_{ss}}{g_{ss}}u'(c_{ss}))dG_{ss}$ with $\int dG_{ss} = 1$ gives the covariance form (45), $\mathrm{Cov}_{G_{ss}}(\omega, \partial_\omega\gamma_{ss}) + \Psi_{ss} = 0$, with the inflation commitment premium $\Psi_{ss} = \int(\omega - A_{ss})\psi_{ss}u'(c_{ss})\,dx$.

Part: Optimal $i^{b,\text{comm}}_{ss}$. Under commitment with $\psi_{ss} \neq 0$, the bond-rate FOC (34) in steady state

with $\phi_{ss} B_{ss} = A^b_{ss}$ becomes

$$\int (\omega - y_{ss})(\partial_\omega \gamma_{ss}\, g_{ss} + \psi_{ss}\, \partial_\omega W_{ss}) dx - \int \eta_{ss}\, u'(c_{ss})\, dx \;=\; \theta_{ss}\, A^b_{ss}.$$

Using the household envelope condition $\partial_\omega W_{ss} = u'(c_{ss})$, then substituting $\theta_{ss} = \int \partial_\omega \gamma_{ss}\, dG_{ss} + \int \psi_{ss} u'(c_{ss})\, dx$ and $A^b_{ss} = \int (\omega - y_{ss})\, dG_{ss}$, $\int dG_{ss} = 1$, this is the covariance form (46), $\mathrm{Cov}_{G_{ss}}(\omega - y_{ss}, \partial_\omega \gamma_{ss}) + \Psi^b_{ss} = \int \eta_{ss} u'(c_{ss})\, dx$, with the bond commitment premium $\Psi^b_{ss} = \int ((\omega - y_{ss}) - A^b_{ss}) \psi_{ss} u'(c_{ss})\, dx$. The discretion condition (41) obtains when $\psi_{ss} \equiv 0$, hence $\Psi^b_{ss} = 0$.

Part: Forward KFE for ψ_{ss}. Setting $\partial_t \psi_s = 0$ in (33) immediately gives the stationary forward KFE for ψ_{ss}. □

A.10 Proof of Lemma 6

The first identity is the definition of the flow consumption wedge under Lemma 5: $\tilde{\chi}_{ss} \equiv \partial_\omega \gamma_{ss} - u'(c_{ss}) = \theta_{ss} \partial_\omega \gamma^f_{ss}$. For the second, write the lumpy consumption wedge (37) on the interior region ($\xi_{ss} = 0$) and substitute the decomposition at both evaluation points, using the envelope condition $\partial_\omega W_{ss} = u'(c_{ss})$:

$$\begin{aligned} \alpha[\partial_\omega \gamma_{ss}(\omega - y_{ss}, z) - U'(y_{ss})] &= \alpha[u'(c_{ss}(\omega - y_{ss}, z)) - U'(y_{ss})] + \alpha\, \theta_{ss}\, \partial_\omega \gamma^f_{ss}(\omega - y_{ss}, z) \\ &= -\, i^b_{ss}\, u'(c_{ss}(\omega, z)) + \alpha\, \theta_{ss}\, \partial_\omega \gamma^f_{ss}(\omega - y_{ss}, z), \end{aligned}$$

where the second line uses the interior lumpy Euler (13) with equality, $\alpha U'(y_{ss}) = \alpha\, \partial_\omega W_{ss}(\omega - y_{ss}, z) + i^b_{ss}\, \partial_\omega W_{ss}(\omega, z)$. Adding the portfolio gap $i^b_{ss}[\partial_\omega \gamma_{ss}(\omega, z) - \theta_{ss}] = i^b_{ss}\, u'(c_{ss}(\omega, z)) + i^b_{ss} \theta_{ss}\, \partial_\omega \gamma^f_{ss}(\omega, z) - i^b_{ss} \theta_{ss}$, the marginal-utility terms cancel and

$$\tilde{\eta}_{ss} \;=\; \theta_{ss} \Big[\alpha\, \partial_\omega \gamma^f_{ss}(\omega - y_{ss}, z) - i^b_{ss}(1 - \partial_\omega \gamma^f_{ss}(\omega, z)) \Big],$$

which is (47). □

A.11 Proof of Lemma 7

Write $h \equiv \partial_\omega \gamma^f_{ss}$ throughout. The proof proceeds in three steps: linearity of the stationary promise multiplier in θ_{ss}, its normalization, and the pairing that delivers (48).

Step 1 (θ-linearity). The stationary version of the forward equation (33) is $\mathcal{A}^*_{ss}[\psi_{ss}] = -\, \partial_\omega \Phi_{ss}$, with Φ_{ss} the flux potential defined in the text. By the multiplier definitions (31) and (32) together with Lemma 6,

$$\chi_{ss} = \theta_{ss}\, \frac{h\, g_{ss}}{u''(c_{ss})}, \qquad \eta_{ss} = \theta_{ss}\, \frac{[\alpha\, h(\omega - y_{ss}, z) - i^b_{ss}(1 - h(\omega, z))]\, g_{ss}}{\alpha[U''(y_{ss}) + \partial_{\omega\omega} W_{ss}(\omega - y_{ss}, z)]},$$

so $\Phi_{ss} = \theta_{ss}\Phi^1_{ss}$ where Φ^1_{ss} contains only household objects and h. The solution set of the stationary equation is therefore $\theta_{ss}\psi^1_{ss} + c\, g_{ss}$, $c \in \mathbb{R}$, since the stationary density spans the null space of $\mathcal{A}^*_{ss}$.

Step 2 (zero mass). Pairing the forward equation with the constant function: $\langle \mathbf{1}, \mathcal{A}^*_{ss}\psi\rangle = \langle \mathcal{A}_{ss}\mathbf{1}, \psi\rangle = 0$, while the right side integrates to the boundary flux of Φ_{ss}, which vanishes—the multipliers carry the density g_{ss}, which decays at the top of the wealth distribution, and the boundary condition from the $\delta W_s(\underline{b}, z)$ variation absorbs the residual at the constraint. Hence $\frac{d}{dt}\int \psi_t\, dx = 0$ along any commitment transition; since the commitment problem starts unencumbered, $\psi_0 = 0$, the stationary limit inherits $\int \psi_{ss}\, dx = 0$, which selects $c = 0$. In particular $\psi_{ss} = \theta_{ss}\psi^1_{ss}$ with $\int \psi^1_{ss}\, dx = 0$.

Step 3 (pairing). Let $q \equiv (\omega - A_{ss})\, u'(c_{ss}) - \mathrm{Cov}_{G_{ss}}(\omega, u'(c_{ss}))$ denote the household's ledger entry, so the Poisson equation defining $\tilde{q}_{ss}$ reads $\mathcal{A}_{ss}[\tilde{q}_{ss}] = -q$, solved by the cumulated entry $\tilde{q}_{ss}(x) = \mathbb{E}_x \int_0^\infty q(x_s)\, ds$. The Poisson equation is well posed: integrating any candidate solution against the stationary distribution gives $\int \mathcal{A}_{ss}[\tilde{q}_{ss}]\, dG_{ss} = 0$, so solvability requires the compatibility condition $\int q\, dG_{ss} = 0$, which q satisfies by construction; existence then follows on the ergodic generator, with $\tilde{q}_{ss}$ unique up to an additive constant because the null space of $\mathcal{A}_{ss}$ is spanned by constants. The constant is annihilated below because $\mathcal{A}_{ss}\mathbf{1} = 0$ and $\int \psi_{ss}\, dx = 0$. Then

$$\Psi_{ss} \;=\; \int (\omega - A_{ss})\, u'(c_{ss})\, \psi_{ss}\, dx \;=\; \int q\, \psi_{ss}\, dx \;=\; -\int \mathcal{A}_{ss}[\tilde{q}_{ss}]\, \psi_{ss}\, dx \;=\; -\int \tilde{q}_{ss}\, \mathcal{A}^*_{ss}[\psi_{ss}]\, dx,$$

where the second equality uses $\int \psi_{ss}\, dx = 0$ to add the constant $-\mathrm{Cov}_{G_{ss}}(\omega, u'(c_{ss}))$ at no cost—this is where the normalization of Step 2 is essential—and the last moves the generator to its adjoint. Substituting the forward equation and integrating by parts in ω with vanishing boundary flux,

$$\Psi_{ss} \;=\; \int \tilde{q}_{ss}\, \partial_\omega \Phi_{ss}\, dx \;=\; -\int \partial_\omega \tilde{q}_{ss}\; \Phi_{ss}\, dx \;=\; -\,\theta_{ss}\int \partial_\omega \tilde{q}_{ss}\; \Phi^1_{ss}\, dx,$$

which is (48). □

A.12 Proof of Proposition 5

Proposition 5 is proved in Appendix D, which derives the separated optimality system—both the commitment ($\lambda \to 0$) and discretion ($\lambda \to \infty$) limits—by the discretize-first method of Dávila and Schaab (2023) in detrended-nominal variables, with the joint continuum limit recovering the conditions stated in Section 4.2.

A.13 Proof of Proposition 6

Under full discretion ($\psi_s \equiv 0$) in steady state, the optimality system of Proposition 5 reduces to the discretion conditions (197) and (198) of Appendix D, with $\theta_{ss} = \int \partial_\omega \gamma_{ss}\, dG_{ss}$:

$$\int \omega\, \partial_\omega \gamma_{ss}\, dG_{ss} = \vartheta_{ss}(A^a_{ss} - A^{b,CB}_{ss}) + \zeta_{ss} A^{b,T}_{ss}, \qquad \int (\omega - y_{ss})\, \partial_\omega \gamma_{ss}\, dG_{ss} - \int \eta_{ss} u'(c_{ss})\, dx = \zeta_{ss} A^{b,T}_{ss} - \vartheta_{ss} A^{b,CB}_{ss}.$$

Covariance form. Subtract $\theta_{ss} \int \omega\, dG_{ss} = \theta_{ss} A_{ss}$ from the first condition, using $A_{ss} = A^a_{ss} + A^{b,T}_{ss} - A^{b,CB}_{ss}$ and $\int \partial_\omega \gamma_{ss}\, dG_{ss} = \theta_{ss}$. The left side becomes $\mathrm{Cov}_{G_{ss}}(\omega, \partial_\omega \gamma_{ss})$ and the right collects into $(\vartheta_{ss} - \theta_{ss})(A^a_{ss} - A^{b,CB}_{ss}) + (\zeta_{ss} - \theta_{ss}) A^{b,T}_{ss}$, which is (76). Subtracting $\theta_{ss} \int (\omega - y_{ss})\, dG_{ss} = \theta_{ss} A^b_{ss}$ (with $A^b_{ss} = A^{b,T}_{ss} - A^{b,CB}_{ss}$) from the second condition gives (77) the same way. □

A.14 Proof of Proposition 7

Under commitment $\psi_{ss} \not\equiv 0$, so the wealth- and bond-weighted social values in the optimal-inflation and bond-rate conditions (59) and (58) keep their promise terms. In steady state, with $\partial_\omega W_{ss} = u'(c_{ss})$ and $\theta_{ss} = \int \partial_\omega \gamma_{ss}\, dG_{ss} + \int \psi_{ss} u'(c_{ss})\, dx$, they read

$$\int \omega\, \partial_\omega \gamma_{ss}\, dG_{ss} + \int \omega\, \psi_{ss} u'(c_{ss})\, dx = \vartheta_{ss}(A^a_{ss} - A^{b,CB}_{ss}) + \zeta_{ss} A^{b,T}_{ss},$$
$$\int (\omega - y_{ss})\, \partial_\omega \gamma_{ss}\, dG_{ss} + \int (\omega - y_{ss})\, \psi_{ss} u'(c_{ss})\, dx = \zeta_{ss} A^{b,T}_{ss} - \vartheta_{ss} A^{b,CB}_{ss} + \int \eta_{ss} u'(c_{ss})\, dx.$$

Subtract $\theta_{ss} A_{ss}$ from the first ($A_{ss} = \int \omega\, dG_{ss} = A^a_{ss} + A^{b,T}_{ss} - A^{b,CB}_{ss}$). The $\partial_\omega \gamma$ terms give $\mathrm{Cov}_{G_{ss}}(\omega, \partial_\omega \gamma_{ss})$, the ψ terms give $\int (\omega - A_{ss})\, \psi_{ss} u'(c_{ss})\, dx = \Psi_{ss}$, and the right collects to $(\vartheta_{ss} - \theta_{ss})(A^a_{ss} - A^{b,CB}_{ss}) + (\zeta_{ss} - \theta_{ss}) A^{b,T}_{ss}$. Subtracting $\theta_{ss} A^b_{ss}$ from the second ($A^b_{ss} = A^{b,T}_{ss} - A^{b,CB}_{ss}$) works the same way, with $\Psi^b_{ss} = \int ((\omega - y_{ss}) - A^b_{ss}) \psi_{ss} u'(c_{ss})\, dx$. The covariance-form conditions are

$$\mathrm{Cov}_{G_{ss}}(\omega,\, \partial_\omega \gamma_{ss}) + \Psi_{ss} = (\vartheta_{ss} - \theta_{ss})(A^a_{ss} - A^{b,CB}_{ss}) + (\zeta_{ss} - \theta_{ss})\, A^{b,T}_{ss}, \tag{74}$$
$$\mathrm{Cov}_{G_{ss}}(\omega - y_{ss},\, \partial_\omega \gamma_{ss}) + \Psi^b_{ss} = \int \eta_{ss}\, u'(c_{ss})\, dx + (\zeta_{ss} - \theta_{ss})\, A^{b,T}_{ss} - (\vartheta_{ss} - \theta_{ss})\, A^{b,CB}_{ss}. \tag{75}$$

Subtracting the first from the second cancels the shared terms $(\zeta_{ss} - \theta_{ss}) A^{b,T}_{ss}$ and $-(\vartheta_{ss} - \theta_{ss}) A^{b,CB}_{ss}$ and gives the combined condition (65) of Proposition 7. Setting $\psi_{ss} \equiv 0$ gives the discretion covariance conditions

$$\mathrm{Cov}_{G_{ss}}(\omega,\, \partial_\omega \gamma_{ss}) = (\zeta_{ss} - \theta_{ss})\, A^{b,T}_{ss} + (\vartheta_{ss} - \theta_{ss})(A^a_{ss} - A^{b,CB}_{ss}), \tag{76}$$
$$\mathrm{Cov}_{G_{ss}}(\omega - y_{ss},\, \partial_\omega \gamma_{ss}) = \int \eta_{ss}\, u'(c_{ss})\, dx + (\zeta_{ss} - \theta_{ss})\, A^{b,T}_{ss} - (\vartheta_{ss} - \theta_{ss})\, A^{b,CB}_{ss}. \tag{77}$$

Subtracting the first from the second cancels $(\zeta_{ss} - \theta_{ss}) A^{b,T}_{ss}$ and $(\vartheta_{ss} - \theta_{ss}) A^{b,CB}_{ss}$ and gives (61), the condition of Proposition 6. □

B Derivation of Kolmogorov Forward Equation and Adjoint Operator

We derive the Kolmogorov forward equation (15) and the adjoint operator $\mathcal{A}_t^*$ by taking limits of a discrete-time economy as the length of a period goes to zero, following closely to Achdou et al. (2022). The derivation is presented for the one-asset (liquid wealth only) case for clarity; the extension to the two-asset model follows analogously with additional drift terms for the illiquid account. I focus on the case of two income states, but the argument extends straightforwardly to any finite, discrete state space for z. Assume time is discrete and the length of a period is Δ. We assume there are only two productivity states, where the transition rate from state $z_t = z_j$ to state $z_{t+\Delta} = z_{-j}$ is $\Delta\lambda_j$. The probability of a lumpy expenditure shock is $\Delta\alpha$. Given beginning of period t wealth $\tilde{a}_t$, households first make their flow savings/consumption decisions, $\Delta s_t(\tilde{a}_t, z_t)$, then learn about their lumpy expenditure shock and choose an amount of lumpy consumption $y_t(\hat{a}_t, z_t)$, where $\hat{a}_t$ is wealth after flow savings/consumption, and then finally households learn about their productivity shock.

A household's beginning of period $t + \Delta$ wealth, $\tilde{a}_{t+\Delta}$, depends on if they receive a lumpy expenditure shock at time t. The law of motion of their wealth is

$$\tilde{a}_{t+\Delta} = \tilde{a}_t + \Delta s_t(\tilde{a}_t, z_t), \text{ or} \tag{78}$$

$$\tilde{a}_{t+\Delta} = \tilde{a}_t + \Delta s_t(\tilde{a}_t, z_t) - y_t(\hat{a}_t, z_t), \tag{79}$$

where $\hat{a}_t \equiv \tilde{a}_t + \Delta s_t(\tilde{a}_t, z_t)$. Notice, there is a important distinction between flow savings, Δs, that depends on the length of a period versus lumpy (dis)savings, y, that does not. We can alternatively write the law of motion (78)-(79) backwards in time as

$$\tilde{a}_t = \tilde{a}_{t+\Delta} - \Delta s_{t+\Delta}(\tilde{a}_{t+\Delta}, z_t), \text{ or} \tag{80}$$

$$\tilde{a}_t = \tilde{a}_{t+\Delta} - \Delta s_{t+\Delta}(\tilde{a}_{t+\Delta}, z_t) + y_t(\hat{a}_{t+\Delta}, z_t). \tag{81}$$

Let $G_t(a, z_j) = Pr(\tilde{a}_t \leq a, \tilde{z}_t = z_j)$ denote the CDF of the wealth distribution of type-j households at time t. We can write

$$\begin{aligned} Pr(\tilde{a}_{t+\Delta} < a, \tilde{z}_{t+\Delta} = z_j) =& (1 - \Delta\lambda_j)(1 - \Delta\alpha) Pr(\tilde{a}_t \leq a - \Delta s_{t+\Delta}(a, z_j), \tilde{z}_t = z_j) \\ &+ + (1 - \Delta\lambda_j)\Delta\alpha Pr(\tilde{a}_t \leq a - \Delta s_{t+\Delta}(a, z_j) + y_{t+\Delta}(a, z_j), \tilde{z}_t = z_j) \\ &+ \Delta\lambda_{-j}(1 - \Delta\alpha) Pr(\tilde{a}_t \leq a - \Delta s_{t+\Delta}(a, z_{-j}), \tilde{z}_t = z_{-j}) \\ &+ \Delta\lambda_{-j}\Delta\alpha Pr(\tilde{a}_t \leq a - \Delta s_{t+\Delta}(a, z_{-j}) + y_{t+\Delta}(a, z_{-j}), \tilde{z}_t = z_{-j}). \end{aligned} \tag{82}$$

In (82), the first term captures the event when the household a type $(\tilde{a}_t, \tilde{z}_t)$ household does not receive either a productivity or lumpy expenditure shock. The second and third terms capture

the events when the household only receives a lumpy expenditure or productivity shock, respectively. The last term captures the event when both shocks occur. Notice, in (82) we implicitly replaced $\hat{a}_{t+\Delta}$ with $\tilde{a}_{t+\Delta}$ in $y_{t+\Delta}$ which is a good approximation when we take limits as $\Delta \to 0$. Using the definition of the CDF, we can write

$$\begin{aligned} G_{t+\Delta}(a, z_j) =& (1 - \Delta\lambda_j)(1 - \Delta\alpha) G_t(a - \Delta s_{t+\Delta}(a, z_j), z_j) \\ &+ + (1 - \Delta\lambda_j)\Delta\alpha G_t(a - \Delta s_{t+\Delta}(a, z_j) + y_{t+\Delta}(a, z_j)) \\ &+ \Delta\lambda_{-j}(1 - \Delta\alpha) G_t(a - \Delta s_{t+\Delta}(a, z_{-j}), z_{-j}) \\ &+ \Delta\lambda_{-j}\Delta\alpha G_t(a - \Delta s_{t+\Delta}(a, z_{-j}) + y_{t+\Delta}(a, z_{-j}), z_{-j}). \end{aligned} \tag{83}$$

Subtracting $G_t(a, z_j)$ from both sides and taking limits as $\Delta \to 0$, we get

$$\begin{aligned} \partial_t G_t(a, z_j) =& -s_t(a, z_j)\partial_a G_t(a, z_j) - \alpha \left[G_t(a, z_j) - G_t(a + y_t(a, z_j), z_j) \right] \\ &- \lambda_j G_t(a, z_j) + \lambda_{-j} G_t(a, z_{-j}). \end{aligned} \tag{84}$$

Differentiating (84) with respect to a, using $g_t(a, z) = \partial_a G_t(a, z)$, yields the KFE (15). The drift term gives $-\partial_a[s_t g_t]$ directly. For the jump term, the chain rule applied to $G_t(a + y_t(a, z), z)$ produces $g_t(a + y_t(a, z), z) \cdot (1 + \partial_a y_t(a, z))$, which is the source of the Jacobian correction $(1 + \partial_a y_t)$ in the adjoint generator. This correction is generically non-zero wherever $y_t(a, z)$ varies with a.

C Proof of Proposition 1

This appendix proves Proposition 1. It re-derives the full optimality system with the discretize-first method of Dávila and Schaab (2023) by first setting up the planner's problem in detrended-nominal variables. The conditions of Proposition 1 are then recovered by taking the joint continuum limit (time step and mesh $\to 0$) of the discretized first-order conditions.

C.1 The detrended nominal economy

We set up and solve the planner's problems — both commitment and discretion — in *detrended-nominal* variables, expressing household wealth and the nominal stocks relative to the money stock M_t. This is a convenient normalization for two reasons. First, nominal quantities trend with M_t, so dividing by it removes the common trend and leads to a well-defined steady state. Second, the detrended variables make each date's problem local — a date's conditions reach into the future only through the forward value (see Lemma 8) — which is particularly useful for the discretion case.

C.1.1 Change of variables

Generally, we denote nominal variables with a tilde ($\tilde{x}$) and detrended nominal variables with a hat ($\hat{x}$). Let $\tilde{\omega} = \omega/\phi_t$ denote a household's total *nominal* wealth and M_t the nominal money stock. Define

$$\hat{\omega}_t \equiv \frac{\tilde{\omega}}{M_t} = \frac{\omega}{\phi_t M_t}, \qquad Q_t \equiv \phi_t M_t \ , \qquad \beta_t \equiv \frac{B_t}{M_t}, \qquad g_t^M \equiv \frac{\dot{M}_t}{M_t}, \tag{85}$$

so that real wealth is $\omega = Q_t\,\hat{\omega}_t$. The aggregate state inherited by a planner is the single distribution $\hat{F}_t$, the measure over $(\hat{\omega}, z)$. Two reductions deliver this. First, the *level* of M_t is payoff-irrelevant by nominal homogeneity (scaling all nominal objects jointly scales ϕ inversely and leaves all real objects unchanged), and detrending by M_t removes it. Second, the bond–money ratio β_t is not independent: asset-market accounting fixes $\int \hat{\omega}\, d\hat{F}_t = 1 + \beta_t$, so $\beta_t = \int \hat{\omega}\, d\hat{F}_t - 1$ is a function of the mean of $\hat{F}_t$. The return on money is a *derived* object,

$$r_t^a = \frac{\dot{\phi}_t}{\phi_t} = \frac{\dot{Q}_t}{Q_t} - g_t^M, \qquad \pi_t = -r_t^a = g_t^M - \frac{\dot{Q}_t}{Q_t}; \tag{86}$$

inflation is money growth net of real-balance growth. At a stationary equilibrium $\dot{Q} = 0$ and $\pi_{ss} = g_{ss}^M$.

C.1.2 Household problem in detrended units

Money pays zero nominal interest and bonds pay nominal i_t^b; real flow income $z + T_t - c_t$ converts to nominal at rate $1/\phi_t$. Nominal wealth therefore accumulates as $\dot{\tilde{\omega}} = i_t^b \tilde{b} + (z + T_t - c_t)/\phi_t$, with nominal money $\tilde{a} = y_t/(\phi_t M_t) \cdot M_t$ chosen so that real lumpy consumption y_t is financed: $\tilde{a} = y_t/\phi_t$, i.e. $\tilde{a}/M_t = y_t/Q_t$. Dividing by M_t and subtracting the dilution term $\hat{\omega}_t\, g_t^M$:

$$\dot{\hat{\omega}} \;=\; \hat{s}_t(\hat{\omega}, z) \;\equiv\; i_t^b\Big(\hat{\omega} - \frac{y_t}{Q_t}\Big) + \frac{z + T_t - c_t}{Q_t} - g_t^M\,\hat{\omega}, \tag{87}$$

Notice we can write $\dot{Q_t\hat{\omega}} = \dot{\omega}$. Let the household's value over the detrended state, $V_t(\hat{\omega}, z) = W_t(Q_t\hat{\omega}, z)$, satisfies

$$\rho V_t = \max_{c,y}\; u(c) + \alpha U(y) + \hat{s}_t\,\partial_{\hat{\omega}} V_t + \alpha\big[V_t(\hat{\omega} - y/Q_t, z) - V_t(\hat{\omega}, z)\big] + \mathcal{A}_z V_t + \partial_t V_t, \tag{88}$$

where the jump in detrended nominal wealth after a lumpy expenditure is $\hat{\omega} - y/Q_t$. The household's first-order conditions are

$$u'(c_t) = \frac{\partial_{\hat{\omega}} V_t}{Q_t}, \qquad \alpha U'(y_t) = \frac{\alpha\,\partial_{\hat{\omega}} V_t(\hat{\omega} - y_t/Q_t, z) + i_t^b\,\partial_{\hat{\omega}} V_t(\hat{\omega}, z)}{Q_t} \quad \text{when } \hat{\omega} \le \underline{\hat{\omega}}_t(z). \tag{89}$$

where $\underline{\hat{\omega}}_t(z) = \underline{\omega}_t(z)/Q_t$ is the full-depletion threshold in detrended nominal wealth.

Borrowing constraint. The real constraint $\omega \ge \underline{b}$ becomes $\hat{\omega} \ge \underline{b}/Q_t$. We consider the case where $\underline{b} = 0$, so the constraint is the *fixed* grid boundary $\hat{\omega} = 0$.[41]

C.1.3 Aggregates, market clearing, government

Real money demand is $\int y\, d\hat{F}_t$ (each household holds real balances equal to intended lumpy consumption), so market clearing in real units is

$$\text{money:}\;\; \int y_t\, d\hat{F}_t = Q_t, \qquad \text{bonds:}\;\; \int (Q_t\hat{\omega} - y_t)\, d\hat{F}_t = Q_t\,\beta_t. \tag{90}$$

Total real public liabilities are $A_t = \phi_t(M_t + B_t) = Q_t(1 + \beta_t)$, with $A_t^a = Q_t$ and $A_t^b = Q_t\beta_t$. The consolidated nominal budget constraint $\dot{M} + \dot{B} = i^b B + (\int T\, d\hat{F})/\phi$, multiplied by $\phi_t = Q_t/M_t$ and written in (g^M, β):

$$Q_t\, g_t^M\,(1 + \beta_t) \;+\; Q_t\,\dot{\beta}_t \;-\; i_t^b\, Q_t\,\beta_t \;-\; \int T_t\, d\hat{F}_t \;=\; 0, \tag{91}$$

using $\phi\dot{M} = Qg^M$ and $\phi\dot{B} = Q(\dot{\beta} + g^M\beta)$.

[41] For $\underline{b} < 0$ the boundary $\hat{\omega} = \underline{b}/Q_t$ moves with the control Q_t, so the δQ_n variation acquires Leibniz terms at the boundary, proportional to $-\underline{b}/Q_n^2$ times boundary-row objects (mass point, flux, and the γ/ψ boundary values); the other variations (g^M, β, T, i^b) do not move the boundary and are unaffected. Deriving these terms is mechanical but should precede quantitative work with $\underline{b} < 0$; alternatively, verify numerically that results are continuous as $\underline{b} \uparrow 0$.

C.1.4 The planner's problem

Controls and states. The planner controls, per slot: household objects $(c, y, V, \hat{F})$ subject to implementability; the *static* price-like controls Q_t (real balances) and i_t^b; the policy flows g_t^M (money growth) and the bond-composition state β_t; and transfers T_t. The map $(Q_t, g_t^M, \beta_t) \leftrightarrow (\phi_t, M_t, B_t)$ is a bijection (given M_0), so the space of controls is *identical* to the real formulation's $(\phi_0, \{r_t^a\}, \{\dot{M}_t\}, \{\dot{B}_t\}, \{T_t\})$ — only the coordinates differ. The aggregate state the planner inherits is the detrended pair $(\hat{F}_0, \beta_0)$.

Objective. Following the stochastic-replanning approach of Dávila and Schaab (2023), a planner in office at date 0 maximizes utilitarian welfare until the next regime change, which arrives at the Poisson time τ_1 with intensity λ, taking $(\hat{F}_0, \beta_0)$ as given:

$$\mathcal{W}_0(\hat{F}_0, \beta_0) = \max \ \mathbb{E}_0\left[\int_0^{\tau_1} e^{-\rho s} \int \mathcal{U}_s \, d\hat{F}_s \, ds \ + \ e^{-\rho\tau_1} \, \mathcal{W}_{\tau_1}(\hat{F}_{\tau_1}, \beta_{\tau_1})\right], \qquad \mathcal{U}_s = u(c_s) + \alpha U(y_s), \tag{92}$$

subject to the decentralized equilibrium imposed as implementability constraints: the household HJB (88); the consumption and lumpy Euler conditions (89), the latter on the no-depletion (interior) region where it holds with equality; the forward Kolmogorov equation $\partial_s \hat{f}_s = \hat{\mathcal{A}}_s^*[\hat{f}_s]$ for the detrended density; the government budget (91); and money- and bond-market clearing (90). Full commitment is $\lambda \to 0$; discretion is $\lambda \to \infty$.

Lagrangian. Attaching the multiplier functions $\hat{\psi}_s$ (HJB), $\hat{\chi}_s$ (consumption Euler), $\hat{\eta}_s$ (lumpy Euler) and $\hat{\gamma}_s$ (Kolmogorov equation), and the scalar multipliers θ_s (budget), μ_s (money clearing) and φ_s (bond clearing), the within-slot Lagrangian is

$$\begin{aligned}
\mathcal{L} = \int_0^{\tau_1} e^{-\rho s} \Big\{ & \int \mathcal{U}_s \, d\hat{F}_s \ + \ \int \hat{\psi}_s[-\rho V_s + \mathcal{U}_s + \partial_s V_s + \hat{\mathcal{A}}_s V_s] \, d\hat{\omega} \, dz \\
& + \int \hat{\chi}_s[u'(c_s) - \frac{1}{Q_s}\partial_{\hat{\omega}} V_s] \, d\hat{\omega} \, dz \\
& + \int \hat{\eta}_s[\alpha U'(y_s) - \frac{1}{Q_s}(\alpha \, \partial_{\hat{\omega}} V_s(\hat{\omega} - y_s/Q_s, z) + i_s^b \, \partial_{\hat{\omega}} V_s)] \, \mathbf{1}_s^{\text{int}} \, d\hat{\omega} \, dz \\
& + \int \hat{\gamma}_s[\hat{\mathcal{A}}_s^* \hat{f}_s - \partial_s \hat{f}_s] \, d\hat{\omega} \, dz \ + \ \theta_s[Q_s g_s^M(1 + \beta_s) + Q_s \dot{\beta}_s - i_s^b Q_s \beta_s - \int T_s \, d\hat{F}_s] \\
& + \ \mu_s[\int y_s \, d\hat{F}_s - Q_s] \ + \ \varphi_s[\int (Q_s \hat{\omega} - y_s) \, d\hat{F}_s - Q_s \beta_s] \Big\} \, ds \ + \ e^{-\rho\tau_1} \mathcal{W}_{\tau_1}(\hat{F}_{\tau_1}, \beta_{\tau_1}).
\end{aligned} \tag{93}$$

The eight terms are social welfare followed by the seven implementability blocks (HJB, consumption Euler, lumpy Euler, Kolmogorov equation, government budget, and the two market-clearing conditions). The three aggregate multipliers $\theta_s, \mu_s, \varphi_s$ deliberately carry no hat: the government budget and the two market-clearing conditions are scalar, real-unit equations that are identical to those of the real-variable formulation so the same scalar constraint carries the same multiplier.

The distributional multipliers $\hat{\psi}_s, \hat{\chi}_s, \hat{\eta}_s, \hat{\gamma}_s$ instead pair with *pointwise* constraints integrated over the state, so they relate to their real-formulation counterparts by the change-of-variables factors $\hat{\psi} = Q\psi$, $\hat{\chi} = Q\chi$, $\hat{\eta} = Q\eta$, and $\hat{\gamma}$ unscaled.[42]

Section C.2 discretizes (93) term by term — time derivatives become forward differences, $\hat{\mathcal{A}}_s$ and its adjoint become the matrices $\mathbf{A}_n, \mathbf{A}_n^*$, and the state integrals $\int (\cdot)\, d\hat{\omega}\, dz$ become the weighted inner product $\langle \cdot, \cdot \rangle$ — yielding the slot-n integrands below.

C.1.5 The locality lemma

The benefit of formulating the planning problem in detrended nominal variables is that the planner's controls at time period n affect future periods only through the state they bequeath — the detrended distribution $\hat{F}_{n+1}$ and the bond–money ratio β_{n+1} — with *no successor-dated price* entering the time period n problem. In the real formulation, the return r_n^a at time period $n$depends on the price ϕ_{n+1}, an extra forward channel that is absent here. The only forward-looking object in period n is then the household continuation value V_{n+1}. This is stated formally in the following lemma. The benefit of this formulation will be clear when considering the problem under discretion.

Lemma 8 (Slot locality)**.** *In the detrended, nominal formulation, every slot-n term of the planner's Lagrangian (welfare, HJB, c-Euler, y-Euler, KFE, government budget, both market-clearing conditions) depends only on slot-n objects* $(c_n, y_n, V_n, \hat{f}_n, Q_n, i_n^b, g_n^M, T_n, \beta_n, \beta_{n+1})$ *and on the forward value* V_{n+1} *(through the HJB time difference). No slot-n term references* Q_m, ϕ_m, *or any other successor-dated price for* $m \geq n+1$.

Proof. By inspection of (87)–(91): the drift, jump, market clearing, and budget all involve $(Q_n, i_n^b, g_n^M, T_n, \beta_\cdot)$ only. The inherited distribution $\hat{f}_n$ is *nominal-detrended* and requires no price to be interpreted: unlike the real formulation, there is no revaluation map $g_n = \phi_n^{-1} f_n(\cdot/\phi_n)$ tying the slot-n density to a price the planner influences. The single forward-looking object is V_{n+1}. □

C.2 Discretized Lagrangian

In this section, we first take a discrete approximation of the detrended nominal economy across time and the state space, take first-order conditions with respect to all the planner's controls, and then take limits as the time and grid intervals go to zero. We compare the resulting first-order conditions to those derived in the main body of the paper in the continuous time, real economy and discuss some subtleties. We discretize the detrended economy in time and across the idiosyncratic state, following the upwind finite-volume scheme of Achdou et al. (2022) and Dávila and Schaab (2023, Appendix C).

[42] The mapping for all terms from the nominal-detrended economy to the real economy are provided in Table 4.

C.2.1 Preliminaries

Time. Discretize time into N steps of length dt, with $t_n = (n-1)dt$. Time derivatives become forward differences, e.g. $\partial_t V_t \to (V_{n+1} - V_n)/dt$ and $\partial_t \hat{F}_t \to (\hat{f}_{n+1} - \hat{f}_n)/dt$.

State space and inner product. Let $j \in \{1, \dots, J_{\hat{\omega}}\}$ index the detrended-wealth grid with values $\hat{\omega}_j$ (potentially non-uniform, e.g. finer near the constraint $\hat{\omega} = 0$) and $k \in \{1, \dots, J_z\}$ the productivity grid with Gauss–Hermite values z_k. The combined index $i \in \{1, \dots, J\}$, $J = J_{\hat{\omega}} J_z$, enumerates grid points $(\hat{\omega}_{j(i)}, z_{k(i)})$. With cell widths $\Delta\hat{\omega}_j$ and Δz_k, set

$$\Delta_i \equiv \Delta\hat{\omega}_{j(i)}\, \Delta z_{k(i)}, \qquad D_\Delta \equiv \operatorname{diag}(\Delta_1, \dots, \Delta_J). \tag{94}$$

Functions on the grid become vectors $h = (h_1, \dots, h_J)^\top$ with $h_i = h(\hat{\omega}_{j(i)}, z_{k(i)})$; for two such vectors h, g the weighted inner product is

$$\langle h, g \rangle \equiv \sum_i h_i\, g_i\, \Delta_i = h^\top D_\Delta\, g, \qquad \int \Psi\, d\hat{F}_n \approx \langle \Psi, \hat{f}_n \rangle, \tag{95}$$

so the slot-n density vector $\hat{f}_n$ represents the detrended distribution $\hat{F}_n$.

Discrete generator. The detrended generator $\hat{\mathcal{A}}_n[h] = \hat{s}_n\, \partial_{\hat{\omega}} h + \alpha[h(\hat{\omega} - y_n/Q_n, z) - h] + \mathcal{A}_z[h]$ of (88) is discretized into $\mathbf{A}_n \in \mathbb{R}^{J \times J}$ piece by piece. Let $D^F_{\hat{\omega}}, D^B_{\hat{\omega}}$ be the forward/backward finite-difference matrices on the $\hat{\omega}$-grid (one-sided at $j = 1, J_{\hat{\omega}}$).

(i) *Drift (upwind).* With $\hat{s}^+_{n,i} = \max(\hat{s}_{n,i}, 0)$ and $\hat{s}^-_{n,i} = \min(\hat{s}_{n,i}, 0)$,

$$\mathbf{S}^{\text{drift}}_n \equiv \operatorname{diag}(\hat{s}^+_n)\, D^F_{\hat{\omega}} + \operatorname{diag}(\hat{s}^-_n)\, D^B_{\hat{\omega}}, \tag{96}$$

where $\hat{s}_n$ is the grid evaluation of the detrended drift (87) (equivalently $\mathbf{S}^{\text{drift}}_n = \operatorname{diag}(\hat{s}_n)\, \mathbf{D}^{\hat{\omega}}_n$ with $\mathbf{D}^{\hat{\omega}}_n$ below).

(ii) *Lumpy jump.* The post-jump state $(\hat{\omega}_i - y_{n,i}/Q_n,\, z_{k(i)})$ generally lies between grid points; the linear-interpolation jump matrix $\hat{\mathbf{J}}_n$ encodes $h(\hat{\omega}_i - y_{n,i}/Q_n, z_{k(i)}) \approx \sum_j (\hat{\mathbf{J}}_n)_{ij}\, h_j$. If $\hat{\omega}_i - y_{n,i}/Q_n \in [\hat{\omega}_{j'}, \hat{\omega}_{j'+1}]$ on the slice $z = z_{k(i)}$, then with $w_i = (\hat{\omega}_{j'+1} - (\hat{\omega}_i - y_{n,i}/Q_n))/(\hat{\omega}_{j'+1} - \hat{\omega}_{j'})$ row i has $(\hat{\mathbf{J}}_n)_{i,(j',k(i))} = w_i$ and $(\hat{\mathbf{J}}_n)_{i,(j'+1,k(i))} = 1 - w_i$ (all other entries zero). The discrete lumpy operator is $\alpha(\hat{\mathbf{J}}_n - \mathbf{I})$.

(iii) *Productivity.* The z-process is a Markov chain with generator $\mathbf{A}_z$ acting only on the z-index (block-diagonal across $\hat{\omega}$-slices, rows summing to zero).

Combining the three pieces,

$$\mathbf{A}_n \equiv \mathbf{S}^{\text{drift}}_n + \alpha(\hat{\mathbf{J}}_n - \mathbf{I}) + \mathbf{A}_z. \tag{97}$$

Upwind derivative. The household Euler equations involve $\partial_{\hat{\omega}} V_n$ evaluated with the same upwind logic as the HJB. Define

$$\mathbf{D}_n^{\hat{\omega}} \equiv \operatorname{diag}(\mathbf{1}\{\hat{s}_n \geq 0\})\, D_{\hat{\omega}}^F + \operatorname{diag}(\mathbf{1}\{\hat{s}_n < 0\})\, D_{\hat{\omega}}^B, \tag{98}$$

the upwind first derivative (the same one-sided choice as $\mathbf{S}_n^{\text{drift}}$ but without the drift prefactor). The post-jump derivative $\partial_{\hat{\omega}} V_n(\hat{\omega} - y_n/Q_n, z)$ in the y-Euler is the composition $\hat{\mathbf{J}}_n \mathbf{D}_n^{\hat{\omega}} V_n$. (Recall from (89) that the household derivatives carry a $1/Q_n$ relative to their real counterparts, $\partial_{\hat{\omega}} V = Q\, \partial_\omega W$.)

Adjoints. The KFE pairs the co-state $\hat{\gamma}_n$ against $\mathbf{A}_n \hat{f}_n$, so we need the adjoint of $\mathbf{A}_n$ in the inner product (95), defined by $\langle h, \mathbf{A}_n g\rangle = \langle \mathbf{A}_n^* h, g\rangle$. For any $\mathbf{B} \in \mathbb{R}^{J\times J}$,

$$\langle h, \mathbf{B} g\rangle = h^\top D_\Delta \mathbf{B}\, g = (D_\Delta^{-1}\mathbf{B}^\top D_\Delta\, h)^\top D_\Delta\, g = \langle D_\Delta^{-1}\mathbf{B}^\top D_\Delta\, h,\ g\rangle,$$

so $\mathbf{B}^* = D_\Delta^{-1}\mathbf{B}^\top D_\Delta$, the D_Δ-conjugated transpose. Applying this to each piece,

$$\mathbf{A}_n^* = D_\Delta^{-1}\mathbf{A}_n^\top D_\Delta = (\mathbf{S}_n^{\text{drift}})^* + \alpha(D_\Delta^{-1}\hat{\mathbf{J}}_n^\top D_\Delta - \mathbf{I}) + \mathbf{A}_z^*. \tag{99}$$

In components $(D_\Delta^{-1}\hat{\mathbf{J}}_n^\top D_\Delta)_{ij} = (\Delta_j/\Delta_i)(\hat{\mathbf{J}}_n)_{ji}$, which measures how density at j feeds the adjoint flow at i through lumpy jumps landing near i; in the continuous limit it becomes the integral operator with the inverse-Jacobian kernel $1/(1 - \partial_{\hat{\omega}} y)$ of the pre-jump map.

Boundary at $\hat{\omega} = 0$. With $\underline{b} = 0$ the constraint is the fixed grid edge $\hat{\omega}_1 = 0$. At the cells $(1, k)$: (a) lumpy consumption $y_{n,(1,k)} = 0$ (the household holds $a = y$ and $a = 0$ at the constraint), so the post-jump state is $\hat{\omega}_1$ itself, row $(1, k)$ of $\hat{\mathbf{J}}_n$ is the identity row, and $\alpha(\hat{\mathbf{J}}_n - \mathbf{I})$ contributes nothing there; (b) the state constraint $\hat{\omega} \geq 0$ forces a *non-negative* drift, $\hat{s}_{n,(1,k)} \geq 0$ for every z_k — a household at the constraint cannot dissave below it (the state-constraint boundary condition of the household value function).

Time summation by parts. For any test sequence $\Psi_n \in \mathbb{R}^J$ and state sequence ξ_n,

$$\sum_{n=s}^{N-1} e^{-\rho t_n}\, dt\, \Psi_n^\top \frac{\xi_{n+1} - \xi_n}{dt} = -\sum_{n=s}^{N-1} e^{-\rho t_n}\, dt\, \frac{\Psi_n - e^{\rho dt}\Psi_{n-1}}{dt}\, \xi_n + \left[e^{-\rho t_N}\Psi_{N-1}^\top \xi_N - e^{-\rho t_s}\Psi_{s-1}^\top \xi_s\right]. \tag{100}$$

This is ordinary summation by parts: with $a_n = e^{-\rho t_n}\Psi_n$ and $b_n = \xi_n$, $\sum_{n=s}^{N-1} a_n(b_{n+1} - b_n) = a_{N-1}b_N - a_s b_s - \sum_{n=s+1}^{N-1}(a_n - a_{n-1})b_n$, and $a_n - a_{n-1} = e^{-\rho t_n}[\Psi_n - e^{\rho dt}\Psi_{n-1}]$. The right-hand side moves a backward difference onto the multipliers and isolates the boundary terms at $n = s$ (the inherited state) and $n = N$.

C.2.2 The discretized Lagrangian

Discretizing the eight terms of the Lagrangian (93) by the rules above (forward differences for ∂_s, the matrices $\mathbf{A}_n, \mathbf{A}_n^*$ for $\hat{\mathcal{A}}_s, \hat{\mathcal{A}}_s^*$, and the inner product $\langle \cdot, \cdot \rangle$ for the state integrals) gives the slot-n integrand $\mathcal{I}_n$ as the sum of:

$$\mathcal{I}_n^{\text{welfare}} = \langle \mathcal{U}_n,\ \hat{f}_n \rangle, \tag{101}$$

$$\mathcal{I}_n^{\text{HJB}} = \left\langle \hat{\psi}_n,\ -\rho V_n + \mathcal{U}_n + \frac{V_{n+1} - V_n}{dt} + \mathbf{A}_n V_n \right\rangle, \tag{102}$$

$$\mathcal{I}_n^{c\text{-Euler}} = \left\langle \hat{\chi}_n,\ u'(c_n) - \frac{1}{Q_n} \mathbf{D}_n^{\hat{\omega}} V_n \right\rangle, \tag{103}$$

$$\mathcal{I}_n^{y\text{-Euler}} = \left\langle \hat{\eta}_n,\ \operatorname{diag}(\mathbf{1}_n^{\text{int}}) \left[\alpha U'(y_n) - \frac{1}{Q_n} (\alpha \hat{\mathbf{J}}_n + i_n^b \mathbf{I}) \mathbf{D}_n^{\hat{\omega}} V_n \right] \right\rangle, \tag{104}$$

$$\mathcal{I}_n^{\text{KFE}} = \left\langle \hat{\gamma}_n,\ \mathbf{A}_n^* \hat{f}_n - \frac{\hat{f}_{n+1} - \hat{f}_n}{dt} \right\rangle, \tag{105}$$

$$\mathcal{I}_n^{\text{govBC}} = \theta_n \left[Q_n g_n^M (1 + \beta_{n+1}) + Q_n \frac{\beta_{n+1} - \beta_n}{dt} - i_n^b Q_n \beta_n - T_n \langle \mathbf{1}^J, \hat{f}_n \rangle \right], \tag{106}$$

$$\mathcal{I}_n^{\text{moneyMC}} = \mu_n [\langle y_n, \hat{f}_n \rangle - Q_n], \tag{107}$$

$$\mathcal{I}_n^{\text{bondMC}} = \varphi_n [\langle Q_n \hat{\boldsymbol{\omega}} - y_n, \hat{f}_n \rangle - Q_n \beta_n], \tag{108}$$

where $\hat{\boldsymbol{\omega}} \in \mathbb{R}^J$ is the grid vector $(\hat{\boldsymbol{\omega}})_i = \hat{\omega}_{j(i)}$ and the drift vector is $\hat{s}_{n,i} = i_n^b(\hat{\omega}_i - y_{n,i}/Q_n) + (z_i + T_n - c_{n,i})/Q_n - g_n^M \hat{\omega}_i$. The discrete government budget (106) comes from dividing the nominal budget $M_{n+1} - M_n + B_{n+1} - B_n = dt\, i_n^b B_n + dt\, T_n/\phi_n$ by M_n, multiplying by Q_n, and writing $B_{n+1}/M_n = \beta_{n+1}(1 + g_n^M dt)$; the cross term yields the $g_n^M \beta_{n+1}$ inside $g_n^M(1 + \beta_{n+1})$.

Translation dictionary. The co-states share a letter with their real counterpart, distinguished by the hat ($\hat{\gamma}, \hat{\psi}, \hat{\chi}, \hat{\eta}$ vs. γ, ψ, χ, η); the value function, density, and distribution instead use distinct letters for detrended vs. real (V/W, $\hat{f}/g$, $\hat{F}/G$). Generic functions are written h. The slot subscript n (and t, s in the continuous-time limit) is a *time* index, orthogonal to the detrended-vs-real distinction. Table 4 translates the detrended objects to the paper's real notation, using $\omega = Q\hat{\omega}$ and $d\omega = Q\, d\hat{\omega}$.

The multiplier scalings follow from requiring the constraint pairings to be invariant: e.g. $\int \hat{\psi}(\hat{\omega})\, [\text{HJB}](\hat{\omega})\, d\hat{\omega} = \int \psi(\omega)\, [\text{HJB}](\omega)\, d\omega$ with the HJB bracket itself invariant (shown below), forces $\hat{\psi} = Q\psi$.

Remark 1. *The detrended $\hat{\gamma}$-HJB is equivalent to the real γ-HJB. For any smooth function $h(\omega)$ with $\hat{h}(\hat{\omega}, t) = h(Q_t \hat{\omega}, t)$: $\partial_t \hat{h}|_{\hat{\omega}} = \partial_t h|_{\omega} + \dot{Q}\, \hat{\omega}\, \partial_\omega h$ and $\hat{s}\, \partial_{\hat{\omega}} \hat{h} = (Q\hat{s})\, \partial_\omega h$ with $Q\hat{s} = s^{real} - \dot{Q}\hat{\omega}$ (from (87) and (86), $s^{real} = r^a \omega + i^b(\omega - y) + z + T - c$). The $\pm \dot{Q}\hat{\omega}\, \partial_\omega h$ terms cancel in the sum $\hat{\mathcal{A}}[\hat{h}] + \partial_t \hat{h} = \mathcal{A}[h] + \partial_t h$, and the jump and z-pieces are invariant ($\hat{\omega} - y/Q \leftrightarrow \omega - y$). So every backward equation derived below*

detrended object	real object	rule
$V(\hat{\omega}, z)$	$W(\omega, z)$	$V = W(Q\hat{\omega})$, $\mathbf{D}^{\hat{\omega}} V \leftrightarrow Q\,\partial_\omega W$
$\hat{\gamma}(\hat{\omega}, z)$	$\gamma(\omega, z)$	$\hat{\gamma} = \gamma(Q\hat{\omega})$, $\mathbf{D}^{\hat{\omega}} \hat{\gamma} \leftrightarrow Q\,\partial_\omega \gamma$
$\hat{f}$ (density in $\hat{\omega}$)	g (density in ω)	$\hat{f}(\hat{\omega}) = Q\,g(Q\hat{\omega})$, $\langle \Psi, \hat{f} \rangle = \int \Psi\, dG$
$\hat{\psi}, \hat{\chi}, \hat{\eta}$ (paired with $d\hat{\omega}$)	ψ, χ, η (paired with $d\omega$)	$\hat{\psi} = Q\psi$, $\hat{\chi} = Q\chi$, $\hat{\eta} = Q\eta$
$\theta, \mu, \varphi, i^b$	same	unchanged (real-unit constraints)
Q, β, g^M	—	$Q = \phi M$, $\beta = B/M$, $g^M = \dot{M}/M$
$A^a = Q$, $A^b = Q\beta$, $A = Q(1+\beta)$	ϕM, ϕB, $\phi(M{+}B)$	via (90)
—	r^a	$r^a \equiv \dot{Q}/Q - g^M$ (derived)

Table 4: Detrended-to-real translation dictionary.

in $\hat{\omega}$-coordinates is pointwise identical *to its real-coordinates counterpart. The same cancellation holds for the forward (adjoint) equations after the density change of variables.*

C.3 Full commitment: $\lambda \to 0$

The slot-0 commitment planner's Lagrangian is $\mathcal{L}^C = \sum_{n=0}^{N-1} e^{-\rho t_n}\, dt\, \mathcal{I}_n$ with the integrands (101)–(108). Following Dávila and Schaab (2023, Appendix B.4), we sum by parts *both* the KFE term (multiplier $\hat{\gamma}$) and the forward-looking HJB term (multiplier $\hat{\psi}$), with predecessor convention $\hat{\psi}_{-1} = 0$. Concretely, the two terms carrying a time difference are the HJB term $\langle \hat{\psi}_n, (V_{n+1} - V_n)/dt \rangle$ and the KFE term $-\langle \hat{\gamma}_n, (\hat{f}_{n+1} - \hat{f}_n)/dt \rangle$. Applying the time-SBP identity (100) to each and, for the KFE, using the adjoint $\langle \hat{\gamma}_n, \mathbf{A}_n^* \hat{f}_n \rangle = \langle \mathbf{A}_n \hat{\gamma}_n, \hat{f}_n \rangle$ to move the generator onto $\hat{\gamma}_n$, the time differences shift onto the multipliers,

$$\langle \hat{\psi}_n, \frac{V_{n+1} - V_n}{dt} \rangle \;\to\; -\Big\langle \frac{\hat{\psi}_n - e^{\rho dt}\hat{\psi}_{n-1}}{dt}, V_n \Big\rangle, \qquad -\langle \hat{\gamma}_n, \frac{\hat{f}_{n+1} - \hat{f}_n}{dt} \rangle \;\to\; +\Big\langle \frac{\hat{\gamma}_n - e^{\rho dt}\hat{\gamma}_{n-1}}{dt}, \hat{f}_n \Big\rangle, \tag{109}$$

and the commitment Lagrangian becomes

$$\mathcal{L}^C = \sum_{n=0}^{N-1} e^{-\rho t_n}\, dt\, \widetilde{\mathcal{I}}_n + e^{-\rho t_0} \langle \hat{\gamma}_0, \hat{f}_0 \rangle - e^{-\rho t_{N-1}} \langle \hat{\gamma}_{N-1}, \hat{f}_N \rangle - e^{-\rho t_0} \langle \hat{\psi}_{-1}, V_0 \rangle + e^{-\rho t_N} \langle \hat{\psi}_{N-1}, V_N \rangle, \tag{110}$$

where $\widetilde{\mathcal{I}}_n$ is (101)–(108) with the two time-difference brackets replaced as above: the HJB term contributes $\langle \hat{\psi}_n, -\rho V_n + \mathcal{U}_n + \mathbf{A}_n V_n \rangle - \langle (\hat{\psi}_n - e^{\rho dt}\hat{\psi}_{n-1})/dt, V_n \rangle$ and the KFE term contributes $\langle \mathbf{A}_n \hat{\gamma}_n + (\hat{\gamma}_n - e^{\rho dt}\hat{\gamma}_{n-1})/dt, \hat{f}_n \rangle$. The lower boundary couples the slot-0 multipliers to the inherited density $\hat{f}_0$ and to the promise $\hat{\psi}_{-1} = 0$; the upper boundary at $n = N$ supplies the terminal conditions. The government budget constraint's own $\dot{\beta} = (\beta_{n+1} - \beta_n)/dt$ difference is *not* summed by parts here — it is handled directly in the $\delta\beta_{n+1}$ variation via two adjacent budget rows. The variations below read the FOCs off $\widetilde{\mathcal{I}}_n$ and these boundary terms. Throughout, "$\to$" denotes the continuum limit — the time step $dt \to 0$ *and* the state-space mesh $(\Delta\hat{\omega}, \Delta z) \to 0$, under which grid vectors

become functions and the discrete operators pass to their continuous counterparts ($\mathbf{D}_n^{\hat{\omega}} \to \partial_{\hat{\omega}}$, $\mathbf{A}_n \to \hat{\mathcal{A}}$, $\langle\cdot,\cdot\rangle \to \int(\cdot)\,d\hat{\omega}\,dz$) — followed by translation to real notation via Table 4. (For the static FOCs, e.g. (114), only the mesh limit and the real translation are at work; the time-difference recursions also use $dt \to 0$.)

C.3.1 Household-block variations: δV_n, δc_n, δy_n, $\delta \hat{f}_{n+1}$

Private valuations, δV_n ($n \geq 1$). V_n enters the post-SBP Lagrangian through the HJB constant pieces $\langle \hat{\psi}_n, -\rho V_n + \mathbf{A}_n V_n \rangle$, the SBP recursion $-[(\hat{\psi}_n - e^{\rho dt}\hat{\psi}_{n-1})/dt]^\top D_\Delta V_n$, and the two Euler equations. Collecting coefficients:

$$\frac{\hat{\psi}_n - e^{\rho dt}\hat{\psi}_{n-1}}{dt} = \mathbf{A}_n^* \hat{\psi}_n - \rho \hat{\psi}_n - \frac{1}{Q_n}(\mathbf{D}_n^{\hat{\omega}})^* \Big[\hat{\chi}_n + (\alpha \hat{\mathbf{J}}_n^* + i_n^b \mathbf{I})\,\mathrm{diag}(\mathbf{1}_n^{\mathrm{int}})\,\hat{\eta}_n\Big], \tag{111}$$

and in the limit, $\partial_t \hat{\psi} = \mathcal{A}^*[\hat{\psi}] - \dfrac{1}{Q}(\partial_{\hat{\omega}})^*[\,\hat{\chi} + (\alpha \hat{J}^* + i^b)\mathbf{1}^{\mathrm{int}}\hat{\eta}\,]$. Translating with $\hat{\psi} = Q\psi$, $\hat{\chi} = Q\chi$, $\hat{\eta} = Q\eta$, $\partial_{\hat{\omega}} = Q\,\partial_\omega$, and the density change of variables, the $\dot{Q}$ cross terms cancel exactly as in Remark 1 (the adjoint drift $-\partial_{\hat{\omega}}(\hat{s}\,\hat{\psi})$ produces $-Q\partial_\omega(s\psi) + \dot{Q}\psi + Q\dot{Q}\hat{\omega}\,\partial_\omega\psi$, and the left side's $\partial_t\hat{\psi}|_{\hat{\omega}} = \dot{Q}\psi + Q\partial_t\psi|_\omega + Q\dot{Q}\hat{\omega}\,\partial_\omega\psi$ absorbs both extra terms), leaving

$$\partial_t \psi \;=\; \mathcal{A}^*[\psi] \;+\; \partial_\omega \chi \;+\; \partial_\omega\left[\frac{\alpha\,\eta(\omega^+, z)}{1 - \partial_\omega y(\omega^+, z)}\right] \;+\; \partial_\omega[i^b\,\eta], \tag{112}$$

identical to the paper's (33). The interior indicator $\mathbf{1}_n^{\mathrm{int}}$ is held fixed in this variation, as it is in the instrument variations below. A perturbation of V_n moves the full-depletion threshold, but the Euler-y residual vanishes at the threshold by (20) and the lumpy policy is continuous there, so the moving boundary contributes only at second order. Its one first-order effect is through the Poisson inflow to the boundary mass point, which the boundary conditions of Section C.5 collect.

Flow consumption, δc_n. Flow consumption enters through welfare $u'(c)\hat{f}$; HJB welfare $\hat{\psi}u'(c)$ and HJB drift $-\hat{\psi}\,\mathbf{D}^{\hat{\omega}}V/Q$ (which cancel pairwise by the household c-Euler (89)); c-Euler $u''(c)\hat{\chi}$; KFE volume drift $-\hat{f}\,\mathbf{D}^{\hat{\omega}}\hat{\gamma}/Q$. Hence

$$\hat{\chi}_n = \mathrm{diag}(u''(c_n))^{-1}\mathrm{diag}(\hat{f}_n)\Big[\frac{1}{Q_n}\mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n - u'(c_n)\Big], \tag{113}$$

with continuum limit

$$\chi = \frac{\partial_\omega \gamma - u'(c)}{u''(c)}\,g, \tag{114}$$

equivalent to the paper's (31).

Lumpy consumption, δy_n. The free variation δy_n is restricted to the no-depletion (interior) region: where y_n is an interior choice, $\delta y_{n,j}$ is free; in the corner $y_{n,j} = \hat{\omega}_j$ is pinned by full depletion, so $\delta y_{n,j} = 0$ there. (Those pinned corner values still enter welfare, the jump, and market clearing, and hence the *other* FOCs — "interior" qualifies only the y-Euler equation derived here.) Restricting to the interior, y_n enters $\widetilde{\mathcal{I}}_n$ through: (i) welfare, $\alpha U(y_n)$; (ii) the HJB term — but the $\hat{\psi}_n$-weighted y-derivative, $\alpha U'(y_n) - \frac{1}{Q_n}(\alpha \hat{\mathbf{J}}_n + i_n^b \mathbf{I})\mathbf{D}_n^{\hat{\omega}} V_n$, vanishes by the household's own lumpy Euler (89), so the HJB contributes nothing; (iii) the y-Euler implementability bracket (multiplier $\hat{\eta}_n$); (iv) the post-SBP KFE term $\langle \mathbf{A}_n \hat{\gamma}_n, \hat{f}_n \rangle$, via the drift $\hat{s}_n$ and the jump $\hat{\mathbf{J}}_n$; and (v) money- and bond-market clearing, linearly.

Two derivative identities drive the algebra. The jump lands at $\hat{\omega} - y_n/Q_n$, so for any y_n-independent f, $\partial(\hat{\mathbf{J}}_n f)_i / \partial y_{n,j} = -\delta_{ij} \frac{1}{Q_n} [\mathcal{D}_n^{\text{post}} f]_i$, where $[\mathcal{D}_n^{\text{post}} f]_i$ is the slope of f's interpolant on the bracket containing the post-jump point (so $[\mathcal{D}_n^{\text{post}} f]_i \to \partial_{\hat{\omega}} f(\hat{\omega}_i - y_{n,i}/Q_n, z_i)$); the $1/Q_n$ is the chain-rule factor from the jump size y/Q. The drift gives $\partial \hat{s}_{n,i} / \partial y_{n,j} = -\delta_{ij}\, i_n^b / Q_n$. Collecting the five channels (dividing out $e^{-\rho t_n} dt\, D_\Delta$), the interior δy_n FOC is, componentwise for $\mathbf{1}_{n,j}^{\text{int}} = 1$,

$$0 = \alpha\, \hat{f}_{n,j}\, U'(y_{n,j}) + \mathcal{H}_{n,j}\, \hat{\eta}_{n,j} - \frac{\hat{f}_{n,j}}{Q_n} [i_n^b (\mathbf{D}_n^{\hat{\omega}} \hat{\gamma}_n)_j + \alpha\, [\mathcal{D}_n^{\text{post}} \hat{\gamma}_n]_j] + (\mu_n - \varphi_n)\, \hat{f}_{n,j}, \tag{115}$$

where $\mathcal{H}_{n,j} = \alpha [U''(y_{n,j}) + \frac{1}{Q_n^2} [\mathcal{D}_n^{\text{post}} \mathbf{D}_n^{\hat{\omega}} V_n]_j]$ is the curvature from channels (i) and (iii). Solving for the multiplier,

$$\hat{\eta}_{n,j} = \frac{\hat{f}_{n,j}\, \tilde{\eta}_{n,j}}{\mathcal{H}_{n,j}}, \qquad \tilde{\eta}_{n,j} = \alpha \Big[\frac{1}{Q_n} [\mathcal{D}_n^{\text{post}} \hat{\gamma}_n]_j - U'(y_{n,j}) \Big] + \frac{i_n^b}{Q_n} (\mathbf{D}_n^{\hat{\omega}} \hat{\gamma}_n)_j + (\varphi_n - \mu_n). \tag{116}$$

Taking the limit — $\frac{1}{Q_n} [\mathcal{D}_n^{\text{post}} \hat{\gamma}_n] \to \partial_\omega \gamma(\omega - y, z)$, $\frac{i_n^b}{Q_n} \mathbf{D}_n^{\hat{\omega}} \hat{\gamma}_n \to i^b\, \partial_\omega \gamma(\omega, z)$, $\frac{1}{Q_n^2} \mathcal{D}_n^{\text{post}} \mathbf{D}_n^{\hat{\omega}} V_n \to \partial_{\omega\omega} W(\omega - y, z)$, and $\hat{\eta} = Q\eta$, $\hat{f} = Qg$ (Table 4) — gives

$$\eta = \frac{g\, \tilde{\eta}}{\alpha [U''(y) + \partial_{\omega\omega} W(\omega - y, z)]}, \qquad \tilde{\eta} = \alpha [\partial_\omega \gamma(\omega - y, z) - U'(y)] + i^b [\partial_\omega \gamma(\omega, z) - \theta], \tag{117}$$

the paper's (32)–(37), where the last step substitutes $\mu - \varphi = i^b \theta$ from Section C.3.2.

Distributional state, $\delta \hat{f}_{n+1}$. After the SBP, the density $\hat{f}_{n+1}$ enters *linearly* through: welfare $\langle \mathcal{U}_{n+1}, \hat{f}_{n+1} \rangle$; the two post-SBP KFE terms $\langle \mathbf{A}_{n+1} \hat{\gamma}_{n+1}, \hat{f}_{n+1} \rangle$ and $\langle (\hat{\gamma}_{n+1} - e^{\rho dt} \hat{\gamma}_n)/dt, \hat{f}_{n+1} \rangle$; the transfer term of (106), $-\theta_{n+1} \langle T_{n+1} \mathbf{1}^J, \hat{f}_{n+1} \rangle$; and the two market-clearing conditions, $\mu_{n+1} \langle y_{n+1}, \hat{f}_{n+1} \rangle$ and $\varphi_{n+1} \langle Q_{n+1} \hat{\boldsymbol{\omega}} - y_{n+1}, \hat{f}_{n+1} \rangle$. Setting $\partial \mathcal{L}^C / \partial \hat{f}_{n+1} = 0$ and dividing out $e^{-\rho t_{n+1}} dt\, D_\Delta$,

$$0 = \mathcal{U}_{n+1} + \mathbf{A}_{n+1} \hat{\gamma}_{n+1} + \frac{\hat{\gamma}_{n+1} - e^{\rho dt} \hat{\gamma}_n}{dt} + \mu_{n+1} y_{n+1} + \varphi_{n+1} (Q_{n+1} \hat{\boldsymbol{\omega}} - y_{n+1}) - \theta_{n+1} T_{n+1} \mathbf{1}^J. \tag{118}$$

As $dt \to 0$ the difference $(\hat{\gamma}_{n+1} - e^{\rho dt}\hat{\gamma}_n)/dt \to \partial_t\hat{\gamma} - \rho\hat{\gamma}$; rearranging (118) (and dropping the now-redundant slot index) gives the backward $\hat{\gamma}$-recursion

$$(\rho - \mathbf{A}_n)\hat{\gamma}_n - \partial_t\hat{\gamma}_n = \mathcal{U}_n + \mu_n y_n + \varphi_n(Q_n\hat{\boldsymbol{\omega}} - y_n) - \theta_n T_n \mathbf{1}^I, \tag{119}$$

which under $Q\hat{\boldsymbol{\omega}} \leftrightarrow \omega$ and Remark 1 is identical to (26).

C.3.2 Aggregate variations: δT_n, δi_n^b, δg_n^M, $\delta\beta_{n+1}$, δQ_n

Transfers, δT_n. Lump-sum transfers, T_n, enter (106) ($-\theta_n\langle \mathbf{1}^I, \hat{f}_n\rangle$) and the drift ($\partial\hat{s}_n/\partial T_n = \mathbf{1}^I/Q_n$, weighted by the KFE volume and the ψ-HJB):

$$\theta_n = \frac{1}{Q_n}\Big[\langle \mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n, \hat{f}_n\rangle + \langle \hat{\psi}_n, \mathbf{D}_n^{\hat{\omega}} V_n\rangle\Big], \tag{120}$$

with continuum limit

$$\theta = \int \partial_\omega\gamma \, dG + \underbrace{\int \psi \, \partial_\omega W \, d\omega}_{\Psi}, \tag{121}$$

equivalent to (28) with commitment premium Ψ.

Nominal bond rate, δi_n^b. The scalar i_n^b enters through the drift, $\partial\hat{s}_{n,i}/\partial i_n^b = \hat{\omega}_i - y_{n,i}/Q_n$, which appears in the $\hat{\psi}_n$-weighted HJB and in the post-SBP KFE volume term $\langle \mathbf{A}_n\hat{\gamma}_n, \hat{f}_n\rangle$; through the y-Euler bracket, with coefficient $-\frac{1}{Q_n}\mathbf{D}_n^{\hat{\omega}}V_n$ on the interior; and through the government budget (106), with coefficient $-\theta_n Q_n\beta_n$. Setting $\partial\mathcal{L}^C/\partial i_n^b = 0$ and dividing out $e^{-\rho t_n}dt$,

$$\langle \mathrm{diag}(\hat{\boldsymbol{\omega}} - y_n/Q_n)\,\mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n,\ \hat{f}_n\rangle + \langle \hat{\psi}_n,\ \mathrm{diag}(\hat{\boldsymbol{\omega}} - y_n/Q_n)\,\mathbf{D}_n^{\hat{\omega}}V_n\rangle - \frac{1}{Q_n}\langle \hat{\eta}_n,\ \mathbf{1}_n^{\mathrm{int}}\,\mathbf{D}_n^{\hat{\omega}}V_n\rangle = \theta_n Q_n\beta_n. \tag{122}$$

In the limit — $\hat{\omega} - y_n/Q_n = (\omega - y)/Q$, $\mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n \to Q\,\partial_\omega\gamma$, $\hat{\psi}_n = Q\psi$, $\mathbf{D}_n^{\hat{\omega}}V_n/Q \to u'(c)$, $\hat{\eta}_n = Q\eta$, and $Q_n\beta_n = A^b$ (Table 4) —

$$\int (\omega - y)\,\partial_\omega\gamma \, dG + \underbrace{\int (\omega - y)\,\psi\,\partial_\omega W \, d\omega}_{\Psi^b} - \int \eta\, u'(c)\,\mathbf{1}^{\mathrm{int}}\, d\omega \;=\; \theta\, A^b, \tag{123}$$

equivalent to (34) and (46) with portfolio premium Ψ^b.

Money growth rate, δg_n^M. The money-growth rate g_n^M appears *only at slot n* (Lemma 8; the level of M is detrended away): in the budget (106) with coefficient $\theta_n Q_n(1 + \beta_{n+1})$, and in the drift with

$\partial \hat{s}_n / \partial g_n^M = -\hat{\boldsymbol{\omega}}$ (dilution), weighted by the KFE volume and the ψ-HJB. The FOC is:

$$\theta_n \, Q_n(1+\beta_{n+1}) = \langle \mathrm{diag}(\hat{\boldsymbol{\omega}}) \, \mathbf{D}_n^{\hat{\omega}} \hat{\gamma}_n, \ \hat{f}_n \rangle + \langle \hat{\psi}_n, \ \mathrm{diag}(\hat{\boldsymbol{\omega}}) \, \mathbf{D}_n^{\hat{\omega}} V_n \rangle. \tag{124}$$

Translating ($\langle \mathrm{diag}(\hat{\boldsymbol{\omega}}) \mathbf{D}^{\hat{\omega}} \hat{\gamma}, \hat{f} \rangle = \int \omega \, \partial_\omega \gamma \, dG$, likewise for the $\hat{\psi}$ term, and $Q(1+\beta) = A$):

$$R_t \ \equiv \ \int \omega \, \partial_\omega \gamma_t \, dG_t + \int \omega \, \psi_t \, \partial_\omega W_t \, d\omega \ = \ \theta_t \, A_t \qquad \text{for all } t. \tag{125}$$

which is equivalent to (35).

Bond-money ratio, $\delta\beta_{n+1}$. The ratio of bonds to money, β_{n+1}, appears in the slot-n government budget (coefficient $\theta_n Q_n(g_n^M + 1/dt)$), the slot-($n$+1) budget (coefficient $\theta_{n+1} Q_{n+1}(-1/dt - i_{n+1}^b)$), and the slot-($n$+1) bond market clearing ($-\varphi_{n+1} Q_{n+1}$). The FOC is:

$$\theta_n Q_n (1 + g_n^M dt) \ = \ e^{-\rho dt} \, Q_{n+1} \Big[\theta_{n+1} (1 + i_{n+1}^b dt) + \varphi_{n+1} \, dt \Big]. \tag{126}$$

Equation (126) links the co-state at adjacent slots; the continuous-time co-state equation follows by expanding each dt-dependent factor to first order. Write $Q_{n+1} = Q_n(1 + \hat{q}_n dt)$ with $\hat{q}_n \equiv (Q_{n+1} - Q_n)/(Q_n dt) \to \dot{Q}/Q$; use $e^{-\rho dt} = 1 - \rho dt + O(dt^2)$ and $\theta_{n+1} = \theta_n + \dot{\theta}_n dt + O(dt^2)$; and, wherever a factor already multiplies an explicit dt, replace $i_{n+1}^b, \varphi_{n+1}, g_n^M$ by their slot-n values. Dividing (126) by Q_n, the prefactor is $e^{-\rho dt}(1 + \hat{q}_n dt) = 1 + (\hat{q}_n - \rho) dt + O(dt^2)$ and the bracket is $\theta_{n+1}(1 + i^b dt) + \varphi \, dt = \theta_{n+1} + (\theta i^b + \varphi) dt + O(dt^2)$, so

$$\theta_n (1 + g^M dt) = \theta_{n+1} + [\theta(\hat{q}_n - \rho) + \theta i^b + \varphi] dt + O(dt^2).$$

The $O(1)$ terms give $\theta_n = \theta_{n+1}$ (the co-state is continuous); the $O(dt)$ terms, with $\theta_{n+1} = \theta_n + \dot{\theta}_n dt$ and dividing by dt, give

$$\dot{\theta}_t \ = \ (\rho + g_t^M - \hat{q}_t - i_t^b)\theta_t - \varphi_t \ = \ (\rho - r_t^a - i_t^b)\, \theta_t - \varphi_t \ = \ (\rho - r_t^b)\, \theta_t - \varphi_t, \tag{127}$$

using the *definition* $r^a \equiv \dot{Q}/Q - g^M$ from (86), which is equivalent to (30).

Real balances, δQ_n. The scalar Q_n enters more channels than any other control: the drift, with $\partial \hat{s}_{n,i} / \partial Q_n = [i_n^b y_{n,i} - (z_i + T_n - c_{n,i})]/Q_n^2$ (through the $\hat{\psi}_n$-HJB and the KFE volume term); the jump, through its location $\hat{\omega} - y_n/Q_n$, with $\partial(\hat{\mathbf{J}}_n f)_i / \partial Q_n = (y_{n,i}/Q_n^2)[\mathcal{D}_n^{\mathrm{post}} f]_i$ (again the HJB and KFE volume); the c- and y-Euler brackets, through their $1/Q_n$ factors (and, for the y-Euler, through $\hat{\mathbf{J}}_n$); both market-clearing conditions; and the government budget. Setting $\partial \mathcal{L}^C / \partial Q_n = 0$ and dividing

out $e^{-\rho t_n} dt$,

$$\begin{aligned} 0 \;=\; & \langle \mathrm{diag}(\partial_{Q_n}\hat{s}_n)\, \mathbf{D}_n^{\hat{\omega}} \hat{\gamma}_n,\ \hat{f}_n \rangle + \langle \hat{\psi}_n,\ \mathrm{diag}(\partial_{Q_n}\hat{s}_n)\, \mathbf{D}_n^{\hat{\omega}} V_n \rangle \\ & + \alpha \langle \mathrm{diag}(y_n/Q_n^2)\, \mathcal{D}_n^{\mathrm{post}} \hat{\gamma}_n,\ \hat{f}_n \rangle + \alpha \langle \hat{\psi}_n,\ \mathrm{diag}(y_n/Q_n^2)\, \mathcal{D}_n^{\mathrm{post}} V_n \rangle \\ & + \frac{1}{Q_n^2} \langle \hat{\chi}_n,\ \mathbf{D}_n^{\hat{\omega}} V_n \rangle + \langle \hat{\eta}_n,\ \mathbf{1}_n^{\mathrm{int}}\, \partial_{Q_n} E_n^y \rangle \\ & + \varphi_n [\langle \hat{\boldsymbol{\omega}}, \hat{f}_n \rangle - \beta_n] - \mu_n + \theta_n [g_n^M (1+\beta_{n+1}) + \frac{\beta_{n+1}-\beta_n}{dt} - i_n^b \beta_n], \end{aligned} \tag{128}$$

where $\partial_{Q_n}\hat{s}_{n,i} = [i_n^b y_{n,i} - (z_i + T_n - c_{n,i})]/Q_n^2$ and $\partial_{Q_n} E_{n,i}^y = \frac{1}{Q_n^2}[(\alpha \hat{\mathbf{J}}_n + i_n^b \mathbf{I}) \mathbf{D}_n^{\hat{\omega}} V_n]_i - \frac{\alpha y_{n,i}}{Q_n^3} [\mathcal{D}_n^{\mathrm{post}} \mathbf{D}_n^{\hat{\omega}} V_n]_i$ is the Q_n-derivative of the y-Euler bracket.

Multiply (128) by Q_n and translate via Table 4. Three of the terms simplify immediately: the household lumpy Euler (89) reduces the $\hat{\eta}_n$ bracket in the third line to $\alpha[U'(y) - y\, \partial_{\omega\omega} W(\omega - y)]$; the market-clearing identity $\langle \hat{\boldsymbol{\omega}}, \hat{f}_n \rangle = 1 + \beta_n$ reduces the market-clearing terms on the last line to $(\varphi_n - \mu_n) A_n^a$; and the government budget term on the last line reduces to $\theta_n T_n$. Taking limits as the time and grid steps to to zero yields:

$$\begin{aligned} 0 \;=\; & \int [i^b y - (z + T - c)] (\partial_\omega \gamma\, g + \psi\, \partial_\omega W)\, d\omega \;+\; \alpha \int y\, (\partial_\omega \gamma(\omega - y)\, g + \psi\, \partial_\omega W(\omega - y))\, d\omega \\ & + \int \chi\, u'(c)\, d\omega \;+\; \alpha \int \eta \left[U'(y) - y\, \partial_{\omega\omega} W(\omega - y) \right] \mathbf{1}^{\mathrm{int}}\, d\omega \;+\; (\varphi_n - \mu_n)\, A_n^a \;+\; \theta_n T_n. \end{aligned} \tag{129}$$

We now reduce (129) directly using other optimality conditions in the continuous time/state limit. The key step is to use the *un-substituted* lumpy wedge

$$\tilde{\eta}_{\mathrm{raw}} = \alpha [\partial_\omega \gamma(\omega - y, z) - U'(y)] + i^b\, \partial_\omega \gamma(\omega, z) - (\mu - \varphi), \tag{130}$$

i.e. (117) *before* the replacement $(\mu - \varphi) \to i^b \theta$. The $(\mu - \varphi)$ here is the δy FOC's market-clearing residual ($+\mu\, g$ from money clearing, $a = y$; $-\varphi\, g$ from bond clearing, $b = \omega - y$); the substituted form $i^b[\partial_\omega \gamma - \theta]$ of (117) would presuppose $\mu - \varphi = i^b \theta$, so $(\mu - \varphi)$ is kept symbolic.

Step 1 (drift bracket). With the identity $i^b y - (z + T - c) = r^b \omega - \hat{s}$, where $\hat{s} = r^a \omega + i^b(\omega - y) + z - c + T$, and the shorthands, $\langle h \rangle_\gamma \equiv \int h\, \partial_\omega \gamma\, g\, d\omega$ and $\langle h \rangle_\psi \equiv \int h\, \psi\, \partial_\omega W\, d\omega$ (for a generic function h), the first term of (129) simplifies to $r^b R - \langle \hat{s} \rangle_\gamma - \langle \hat{s} \rangle_\psi$, where $R = \langle \omega \rangle_\gamma + \langle \omega \rangle_\psi$ and, by the transfer FOC (121), $\theta = \langle 1 \rangle_\gamma + \langle 1 \rangle_\psi$. Group (129) by channel. With the Step-1 identity $i^b y - (z + T - c) = r^b \omega - \hat{s}$ and $\hat{s} = s^\omega$ the real drift, the first four terms can be written as

$$\begin{aligned} \text{drift:} \quad & \langle r^b \omega - \hat{s} \rangle_\gamma + \langle r^b \omega - \hat{s} \rangle_\psi = r^b R - \langle \hat{s} \rangle_\gamma - \langle \hat{s} \rangle_\psi, \\ \text{jump:} \quad & \alpha \int y\, \partial_\omega \gamma(\omega - y)\, dG + \alpha \int y\, \partial_\omega W(\omega - y)\, \psi\, d\omega, \\ \chi: \quad & \int \chi\, u'(c)\, d\omega = \int \chi\, \partial_\omega W\, d\omega, \qquad \eta: \quad \alpha \int \eta [U'(y) - y\, \partial_{\omega\omega} W(\omega - y)] \mathbf{1}^{\mathrm{int}}\, d\omega, \end{aligned} \tag{131}$$

using $u'(c) = \partial_\omega W$ throughout. We reduce the γ- and ψ-*channels* that show up in the drift and jump terms in parallel using two results presented in the following lemmas, then assemble them with the χ- and η-blocks.

Lemma 9 (Pointwise commutation)**.** *For any smooth f, the wealth generator $\mathcal{A}h = s^\omega \partial_\omega h + \alpha\,[\,h(\omega - y, z) - h\,] + \mathcal{A}_z h$ satisfies the pointwise identity*

$$\omega\,\mathcal{A}[f] \;=\; \mathcal{A}[\omega f] \;+\; \alpha\, y\, f(\omega - y, z) \;-\; s^\omega f. \tag{132}$$

Proof. The z-generator and the $-\alpha h$ term are unweighted, so they drop out of the difference: $\omega\mathcal{A}[f] - \mathcal{A}[\omega f] = \omega s^\omega \partial_\omega f - s^\omega(f + \omega\partial_\omega f) + \alpha\omega f(\omega - y) - \alpha(\omega - y)f(\omega - y) = -s^\omega f + \alpha y\, f(\omega - y)$. □

Lemma 10 (Flow-derivative identity)**.** *Under the household c-Euler and interior lumpy Euler conditions, the flow utility $\mathcal{U} = u(c) + \alpha U(y)$ satisfies*

$$\partial_\omega \mathcal{U} = (r^b - \partial_\omega s^\omega)\,\partial_\omega W + \alpha\,\partial_\omega y\,\partial_\omega W(\omega - y). \tag{133}$$

Proof. Differentiating $\mathcal{U} = u(c) + \alpha U(y)$ and inserting the c-Euler $u'(c) = \partial_\omega W$ and the interior lumpy Euler $\alpha U'(y) = \alpha\,\partial_\omega W(\omega - y) + i^b\,u'(c)$,

$$\partial_\omega \mathcal{U} = u'(c)\partial_\omega c + \alpha U'(y)\partial_\omega y = \partial_\omega W\,(\partial_\omega c + i^b \partial_\omega y) + \alpha\,\partial_\omega y\,\partial_\omega W(\omega - y). \tag{134}$$

With $\partial_\omega s^\omega = r^b - i^b\partial_\omega y - \partial_\omega c$ (from $s^\omega = r^a\omega + i^b(\omega - y) + z + T - c$, $r^b = r^a + i^b$) we have $\partial_\omega c + i^b\partial_\omega y = r^b - \partial_\omega s^\omega$, which turns (134) into (133). □

Step 2 (γ-channel). Differentiate the γ-HJB (119) (in real coordinates by Remark 1) in ω. Writing the clearing term as $\mu y + \varphi(\omega - y) = \varphi\omega + (\mu - \varphi)y$ so its ω-derivative is $\varphi + (\mu - \varphi)\partial_\omega y$ (the constant $-\theta T$ drops): $\rho\,\partial_\omega\gamma - \partial_t(\partial_\omega\gamma) = \partial_\omega\mathcal{U} + \varphi + (\mu - \varphi)\partial_\omega y + \partial_\omega(\mathcal{A}\gamma)$. Weighting by ω and integrating against dG,

$$\int \omega\,\partial_\omega(\mathcal{A}\gamma)\,dG = \rho\langle\omega\rangle_\gamma - \int \omega\,\partial_t(\partial_\omega\gamma)\,dG - \int \omega\,\partial_\omega\mathcal{U}\,dG - \varphi A - (\mu - \varphi)\int \omega\,\partial_\omega y\,dG, \tag{135}$$

using $\int \omega\,dG = A$ and $\langle\omega\rangle_\gamma = \int \omega\,\partial_\omega\gamma\,dG$. To connect (135) to the drift/jump blocks of (131), apply the commutation identity $\partial_\omega(\mathcal{A}\gamma) = \mathcal{A}[\partial_\omega\gamma] + (\partial_\omega s^\omega)\partial_\omega\gamma - \alpha\,\partial_\omega y\,\partial_\omega\gamma(\omega - y)$, then Lemma 9 with $f = \partial_\omega\gamma$ under $\int(\cdot)\,dG$. The forward KFE $\mathcal{A}^* g = \partial_t g$ gives $\int \mathcal{A}[\omega\,\partial_\omega\gamma]\,dG = \int \omega\,\partial_\omega\gamma\,\partial_t g\,d\omega$, so $\int \omega\,\mathcal{A}[\partial_\omega\gamma]\,dG = \alpha\int y\,\partial_\omega\gamma(\omega - y)\,dG - \int s^\omega\partial_\omega\gamma\,dG + \int \omega\,\partial_\omega\gamma\,\partial_t g\,d\omega = (\text{jump}_\gamma) - \langle\hat{s}\rangle_\gamma + \int \omega\,\partial_\omega\gamma\,\partial_t g\,d\omega$. Hence the γ-drift-plus-jump block $r^b\langle\omega\rangle_\gamma - \langle\hat{s}\rangle_\gamma + (\text{jump}_\gamma)$ equals $r^b\langle\omega\rangle_\gamma + \int \omega\,\mathcal{A}[\partial_\omega\gamma]\,dG -$

$\int \omega\, \partial_\omega \gamma\, \partial_t g\, d\omega$, and substituting (135) for $\int \omega\, \partial_\omega(\mathcal{A}\gamma)\, dG$ and the commutation expansion gives

$$\begin{aligned}(\text{drift}_\gamma + \text{jump}_\gamma) = r^b \langle\omega\rangle_\gamma + \rho\langle\omega\rangle_\gamma - \partial_t\langle\omega\rangle_\gamma - \int \omega\, \partial_\omega \mathcal{U}\, dG - \varphi A - (\mu - \varphi)\int \omega\, \partial_\omega y\, dG \\ - \int \omega(\partial_\omega s^\omega)\partial_\omega \gamma\, dG + \alpha \int \omega\, \partial_\omega y\, \partial_\omega \gamma(\omega - y)\, dG,\end{aligned} \tag{136}$$

where the time derivative collects the two pieces just isolated, $\partial_t\langle\omega\rangle_\gamma = \int \omega[\partial_t(\partial_\omega\gamma)\, g + \partial_\omega\gamma\, \partial_t g]\, d\omega$ (the γ-HJB term and the KFE term).

Step 3 (ψ-channel). Lemma 9 again, now with $f = \partial_\omega W$, integrated against the signed measure $\psi\, d\omega$, gives $\int \omega\, \mathcal{A}[\partial_\omega W]\, \psi\, d\omega = \int (\omega\, \partial_\omega W)\, \mathcal{A}^*[\psi]\, d\omega + \alpha \int y\, \partial_\omega W(\omega - y)\, \psi\, d\omega\, d\omega - \langle \hat{s} \rangle_\psi$, where the first term used the adjoint $\int (\mathcal{A}h)\psi\, d\omega = \int h\, \mathcal{A}^*[\psi]\, d\omega$. Thus the ψ-drift-plus-jump block is

$$(\text{drift}_\psi + \text{jump}_\psi) = r^b \langle\omega\rangle_\psi + \int \omega\, \mathcal{A}[\partial_\omega W]\, \psi\, d\omega - \int \omega\, \partial_\omega W\, \mathcal{A}^*[\psi]\, d\omega. \tag{137}$$

For the second term in (137), differentiate the *private* HJB $\rho W - \partial_t W = \mathcal{U} + \mathcal{A}W$ in ω and apply the commutation identity: $\mathcal{A}[\partial_\omega W] = \rho\, \partial_\omega W - \partial_t(\partial_\omega W) - \partial_\omega \mathcal{U} - (\partial_\omega s^\omega)\partial_\omega W + \alpha\, \partial_\omega y\, \partial_\omega W(\omega - y)$. Substituting the flow identity (Lemma 10) for $\partial_\omega \mathcal{U}$, the $(\partial_\omega s^\omega)\partial_\omega W$ and post-jump terms cancel *exactly*, leaving $\mathcal{A}[\partial_\omega W] = (\rho - r^b)\, \partial_\omega W - \partial_t(\partial_\omega W)$; hence $\int \omega\, \mathcal{A}[\partial_\omega W]\, \psi\, d\omega = (\rho - r^b)\langle\omega\rangle_\psi - \int \omega\, \partial_t(\partial_\omega W)\, \psi\, d\omega$. For the third term in (137), the ψ-KFE (112) reads $\mathcal{A}^*[\psi] = \partial_t \psi - \partial_\omega \chi - \partial_\omega G_L - \partial_\omega[i^b \eta]$ with $G_L \equiv \alpha\, \eta(\omega^+, z)/(1 - \partial_\omega y(\omega^+, z))$, so $-\int \omega\, \partial_\omega W\, \mathcal{A}^*[\psi]\, d\omega = \text{SRC} - \int \omega\, \partial_\omega W\, \partial_t \psi\, d\omega$, where $\text{SRC} \equiv \int \omega\, \partial_\omega W\, (\partial_\omega \chi + \partial_\omega G_L + \partial_\omega[i^b \eta])\, d\omega$. Therefore

$$(\text{drift}_\psi + \text{jump}_\psi) = \rho\langle\omega\rangle_\psi - \partial_t\langle\omega\rangle_\psi + \text{SRC}, \tag{138}$$

where, as in the γ-channel, the two time-derivative pieces collect into $\partial_t\langle\omega\rangle_\psi = \int \omega[\partial_t \psi\, \partial_\omega W + \psi\, \partial_t(\partial_\omega W)]\, d\omega$ (the ψ-KFE term and the private-HJB term). The ψ-flow/drift/post-jump structure collapses to $\rho\langle\omega\rangle_\psi - \partial_t\langle\omega\rangle_\psi$; only SRC remains, to be matched against χ, η in Step 4.

Step 4 (assembly, pairwise cancellation). Collecting (136), (138), the χ, η blocks of (131), and the terms $(\varphi - \mu)A^a + \theta T$:

$$\begin{aligned}0 = r^b \langle\omega\rangle_\gamma + \underbrace{\rho\langle\omega\rangle_\gamma + \rho\langle\omega\rangle_\psi}_{=\rho R} - \underbrace{(\partial_t\langle\omega\rangle_\gamma + \partial_t\langle\omega\rangle_\psi)}_{=\dot{R}} - \varphi A - (\mu - \varphi)\int \omega\, \partial_\omega y\, dG + (\varphi - \mu)A^a + \theta T \\ - \int \omega\, \partial_\omega \mathcal{U}\, dG - \int \omega(\partial_\omega s^\omega)\partial_\omega \gamma\, dG + \alpha \int \omega\, \partial_\omega y\, \partial_\omega \gamma(\omega - y)\, dG + \text{SRC} + \chi\text{-block} + \eta\text{-block}.\end{aligned} \tag{139}$$

(i) *Clearing.* With $A = A^a + A^b$, the clearing line is $-\varphi A^b - \mu A^a + \theta T - (\mu - \varphi)\int \omega\, \partial_\omega y\, dG = \Omega - (\mu - \varphi)\int \omega\, \partial_\omega y\, dG$, where $\Omega \equiv \theta T - \mu A^a - \varphi A^b$.

(ii) *The γ-flow term* becomes a pure wedge. Substituting the flow identity (Lemma 10) for $\partial_\omega \mathcal{U}$

in $-\int \omega\, \partial_\omega \mathcal{U}\, dG$ and combining with the γ drift-derivative and post-jump terms of (139) produces, with $\tilde{\chi} \equiv \partial_\omega \gamma - \partial_\omega W$,

$$\int \omega (r^b - \partial_\omega s^\omega)\tilde{\chi}\, dG + \alpha \int \omega\, \partial_\omega y\, \tilde{\chi}(\omega - y)\, dG \;=\; \int \omega (i^b \partial_\omega y + \partial_\omega c)\tilde{\chi}\, dG + \alpha \int \omega\, \partial_\omega y\, \tilde{\chi}(\omega - y)\, dG, \tag{140}$$

formed as follows: the γ-channel's $r^b \langle\omega\rangle_\gamma$ and the $-r^b \int \omega\, \partial_\omega W\, dG$ from (133) combine into $r^b \int \omega\, \tilde{\chi}\, dG$; the commutation term $-\int \omega(\partial_\omega s^\omega)\partial_\omega \gamma\, dG$ and the $+\int \omega(\partial_\omega s^\omega)\partial_\omega W\, dG$ from (133) combine into $-\int \omega(\partial_\omega s^\omega)\tilde{\chi}\, dG$; and the post-jump γ-term pairs with the post-jump flow term into the $\tilde{\chi}(\omega - y)$ integral. The final equality uses $r^b - \partial_\omega s^\omega = i^b \partial_\omega y + \partial_\omega c$.

(iii) SRC *against* χ, η. Integrate SRC by parts, moving ∂_ω onto $\omega\, \partial_\omega W$ (so $\partial_\omega(\omega\, \partial_\omega W) = \partial_\omega W + \omega\, \partial_{\omega\omega} W$). The leading $\partial_\omega W$ pieces cancel the χ- and η-blocks of (131) termwise: the $-\int \chi\, \partial_\omega W$ from SRC_χ cancels $\int \chi\, \partial_\omega W$; the $-i^b \int \eta\, \partial_\omega W$ from SRC_{i^b} cancels the $i^b\, u'(c)$ piece of the η-block (after $\alpha U'(y) = \alpha\, \partial_\omega W(\omega - y) + i^b u'(c)$); and a change of variables $\omega = \omega^+ - y(\omega^+)$ in SRC_L (Jacobian $1 - \partial_\omega y$ cancels the G_L denominator) turns its leading piece into $-\alpha \int \eta\, \partial_\omega W(\omega - y)$, cancelling the $\alpha U'(y) \to \partial_\omega W(\omega - y)$ piece of the η-block. What survives is the $\omega\, \partial_{\omega\omega} W$ remainder

$$-\int \omega\, \chi\, \partial_{\omega\omega} W\, d\omega - i^b \int \omega\, \eta\, \partial_{\omega\omega} W\, d\omega - \alpha \int \omega\, \eta\, \partial_{\omega\omega} W(\omega - y)\, d\omega, \tag{141}$$

where in the SRC_L remainder the weights $(\omega - y)$ and the bracket's $-y$ post-jump term combine to ω (so $[(\omega - y) + y] = \omega$). Now insert the wedge definitions (114) $\chi = \tilde{\chi}\, g / u''(c)$ and (116) $\eta\, \mathcal{H} = g\, \tilde{\eta}$ with $\mathcal{H} = \alpha[U''(y) + \partial_{\omega\omega} W(\omega - y)]$, together with the differentiated Euler conditions $\partial_\omega c = \partial_{\omega\omega} W / u''(c)$ and $\mathcal{H}\, \partial_\omega y = \alpha\, \partial_{\omega\omega} W(\omega - y) + i^b\, \partial_{\omega\omega} W$ (differentiating the interior lumpy Euler). The first term of (141) becomes $-\int \omega\, \partial_\omega c\, \tilde{\chi}\, dG$, cancelling the $\partial_\omega c$ part of (140); the last two combine to $-\int \omega\, \partial_\omega y\, \tilde{\eta}\, dG$. Writing the *un-substituted* raw wedge (130) in terms of $\tilde{\chi}$ via the lumpy Euler, $\tilde{\eta} = \alpha\, \tilde{\chi}(\omega - y) + i^b\, \tilde{\chi} - (\mu - \varphi)$, this is

$$-\alpha \int \omega\, \partial_\omega y\, \tilde{\chi}(\omega - y)\, dG - i^b \int \omega\, \partial_\omega y\, \tilde{\chi}\, dG + (\mu - \varphi) \int \omega\, \partial_\omega y\, dG, \tag{142}$$

whose first two integrals exactly cancel the $i^b \partial_\omega y$ and post-jump pieces remaining in (140).

The net of (ii)+(iii) is therefore $+(\mu - \varphi) \int \omega\, \partial_\omega y\, dG$, which cancels the leftover $-(\mu - \varphi) \int \omega\, \partial_\omega y\, dG$ from (i). Every flow, post-jump, χ, and η term has cancelled pairwise; the two surviving time derivatives combine into $\dot{R} = \partial_t(\langle\omega\rangle_\gamma + \langle\omega\rangle_\psi)$, and (129) collapses to the co-state ODE

$$\dot{R} = \rho R + \Omega, \qquad \Omega \equiv \theta T - \mu A^a - \varphi A^b, \tag{143}$$

holding at every t. This is the detrended price-path first-order condition; in forward-integral

form, $R_s = -\int_s^\infty e^{-\rho(\tau-s)}\Omega_\tau\, d\tau$ under the transversality condition $e^{-\rho\tau}R_\tau \to 0$.

Boundary terms. The integrations by parts in Step 3 ($\int(\mathcal{A}h)\psi = \int h\,\mathcal{A}^*\psi$, and the SRC reduction) generate boundary contributions at $\omega \to \infty$ and at the lower bound $\omega = \underline{b}$. At $\omega \to \infty$ they vanish under the finite-moment decay of g and ψ (and of $\chi, \eta \sim g$). At $\omega = \underline{b} = 0$ every surviving boundary term carries an explicit weight ω (from the ω-weighting and from $\partial_\omega(\omega\,\partial_\omega W)$), so the factor $\omega = \underline{b} = 0$ annihilates it, and the adjoint step uses zero-flux/mass conservation at the lower bound. For $\underline{b} < 0$ the explicit weight is $\underline{b} \neq 0$ and a genuine boundary term $-\underline{b}[\partial_\omega W(\chi + G_L + i^b\eta)]_{\omega=\underline{b}}$ survives; this moving-boundary Leibniz term — noted at the borrowing constraint above — is set aside under the maintained $\underline{b} = 0$.

Step 5. Write $\Omega = \theta T - \varphi A - (\mu - \varphi)A^a$ and substitute the government budget $T = \dot{A} - r^a A - i^b A^b$ together with the bond co-state $\varphi = (\rho - r^b)\theta - \dot{\theta}$ from (127). Using $r^b = r^a + i^b$ and $A^a = A - A^b$ one finds $\theta T - \varphi A = \dfrac{d}{dt}(\theta A) - \rho\theta A + i^b\theta A^a$, so (143) becomes

$$\frac{d}{dt}(R - \theta A) \;=\; \rho\,(R - \theta A) \;-\; A^a\,(\mu - \varphi - i^b\theta)\,. \tag{144}$$

The coefficient on $(\mu - \varphi - i^b\theta)$ is $-A^a < 0$, so imposing $R = \theta A$ does not render (144) vacuous: δQ carries information independent of δg^M. The static revaluation FOC (125) ($R = \theta A$ at every t) annihilates both $R - \theta A$ and its time derivative, leaving the relation

$$\mu_t - \varphi_t = i^b_t\,\theta_t, \tag{145}$$

equivalent to (36).

Initial conditions. The SBP boundary terms at $n = -1$ carry the inherited promise, $\hat{\psi}_{-1} = 0$. There is no separate ϕ_0 FOC: the date-0 revaluation margin is the $\delta g_0^M/\delta Q_0$ pair, because *every* slot's price variation in this formulation already revalues the (nominal) inherited state.

C.3.3 Summary: the commitment systems coincide

Proposition 8 (Commitment cross-check)**.** *In the continuum limit (time step and mesh $\to$ 0), the detrended-nominal commitment optimality system is, condition for condition, the commitment system used in the main body of the paper:*

detrended variation	*discretized FOC*	*continuum limit*	*paper counterpart*
δc	(113)	(114)	(31)
δy	(115)	(117)	(32)
δV	(111)	(112)	(33)
$\delta \hat{F}$	(118)	(119)	(26)
δT	(120)	(121)	(28) *(with* Ψ*)*
δi^b	(122)	(123)	(34) *(with* Ψ^b*)*
δg^M	(124)	(125)	(35) *(*$\forall t$*)*
$\delta \beta$	(126)	(127)	(30)
δQ	(128)	(144)	(29)

C.4 Discretion: $\lambda \to \infty$

We now characterize the problem of a planner under full discretion as a *Markov perfect equilibrium* (MPE) in the discretized economy: The slot-s planner maximizes welfare from s onward, taking the successor's value $\mathcal{W}_{s+1}(\hat{F}_{s+1}, \beta_{s+1})$ as a fixed function of the state it bequeaths — it shapes the continuation only through the inherited $(\hat{F}_{s+1}, \beta_{s+1})$. Since $\mathcal{W}_{s+1}$ is the successor's own optimum, best-responding slot by slot makes the plan subgame-perfect. The $dt \to 0$ limit of this slot game is the continuous-time $\lambda \to \infty$ system.

The slot-s discretion planner's Lagrangian is $\mathcal{L}^D = \sum_{n=s}^{N-1} e^{-\rho t_n}\, dt\, \mathcal{I}_n$ plus the continuation value $e^{-\rho t_{s+1}} \mathcal{W}_{s+1}(\hat{F}_{s+1}, \beta_{s+1})$, with the integrands (101)–(108). The one operational difference from commitment is that, following Dávila and Schaab (2023), we sum by parts *only* the KFE and leave the HJB term raw: the commitment time-SBP carries the promise co-state ψ backward across slots, whereas discretion does not, so ψ cannot transmit promises and the δV_s variation forces $\psi_s \to 0$ (148) — the discrete counterpart of the $\lambda \to \infty$ condition $\psi \equiv 0$. The structure of most of the variations exactly mirrors the commitment case, so we record the steps compactly and dwell only on what is new.

Private values, δV_s. V_s enters the slot-s Lagrangian in two places: the HJB term, where it appears inside the multiplier product $\psi_s^\top [\, \frac{1}{dt} + \rho - \mathbf{A}_s \,] V_s$ (the discount, the generator, and the *unsummed* time difference carrying the $\frac{1}{dt}$), and the two Euler brackets, $-\frac{1}{Q_s} \langle \mathbf{D}_s^{\hat{\omega}} V_s,\ \chi_s + (\alpha \hat{\mathbf{J}}_s^* + i_s^b \mathbf{I}) \mathrm{diag}(\mathbf{1}_s^{\mathrm{int}}) \eta_s \rangle$. Differentiating in V_s (and moving each operator onto its adjoint) sets $\partial \mathcal{L}^D / \partial V_s = \mathbf{0}$:

$$\left[\mathbf{A}_s^* - \rho - \frac{1}{dt} \right] \psi_s = \frac{1}{Q_s} (\mathbf{D}_s^{\hat{\omega}})^* [\chi_s + (\alpha \hat{\mathbf{J}}_s^* + i_s^b \mathbf{I}) \mathrm{diag}(\mathbf{1}_s^{\mathrm{int}}) \eta_s] \ \equiv\ \mathbf{b}_s. \tag{146}$$

The source $\mathbf{b}_s$ is built from the slot-s multipliers χ_s, η_s and spatial operators only, so it carries no factor of $1/dt$; on the left, because the HJB was *not* summed by parts in time, the operator

on ψ_s is dominated by the time term $-\frac{1}{dt}$ (the $\mathbf{A}_s^* - \rho$ block does not scale with $1/dt$). Writing $\mathbf{A}_s^* - \rho - \frac{1}{dt} = -\frac{1}{dt}[\mathbf{I} - dt(\mathbf{A}_s^* - \rho)]$ and inverting,

$$\psi_s = -\,dt\,[\mathbf{I} - dt(\mathbf{A}_s^* - \rho)]^{-1}\mathbf{b}_s = -\,dt\,\mathbf{b}_s + O(dt^2), \tag{147}$$

so ψ_s is first order in the slot length and vanishes in the continuum limit,

$$\psi_s = O(dt) \;\longrightarrow\; \mathbf{0}. \tag{148}$$

Under commitment the time-SBP instead couples slot s to $s-1$, replacing the bare $-\frac{1}{dt}\psi_s$ by the finite difference $\frac{1}{dt}(\psi_s - e^{\rho dt}\psi_{s-1}) \to \partial_t\psi - \rho\psi$ — the recursion (111) that *propagates* ψ across slots rather than pinning it to zero.

The remaining variations With $\psi_s = \mathbf{0}$, the variations δc_s, δy_s, $\delta \hat{f}_{s+1}$, δT_s, δi_s^b, δg_s^M, δQ_s are *identical* to Section C.3 with the ψ-channels deleted, because each is supported on slot s only (Lemma 8) and the commitment/discretion distinction enters solely through which terms are summed by parts — which affects only the V- and time-boundary structure, not the slot-local channels.

The variation $\delta\beta_{s+1}$ deserves comment: β_{s+1} is a *state* handed to the successor, so the slot-$(s{+}1)$ terms in (126) are replaced by the continuation derivative $\partial\mathcal{W}_{s+1}/\partial\beta_{s+1}$. By the successor's envelope condition this derivative equals the same multiplier combination $e^{-\rho dt}Q_{s+1}[\theta_{s+1}(1 + i^b dt) + \varphi_{s+1}dt]$ evaluated at the successor's optimum (the MPE consistency condition, exactly parallel to the $\delta M_{s+1}, \delta B_{s+1}$ treatment behind the paper's terminal condition for θ_{τ_1}), so the limit recursion is unchanged. The system is:

$$\chi = \frac{\partial_\omega\gamma - u'(c)}{u''(c)}\,g, \qquad \eta = \frac{g\,\tilde{\eta}}{\alpha[U'' + \partial_{\omega\omega}W(\omega - y)]}, \quad \tilde{\eta} = \alpha[\partial_\omega\gamma(\omega - y) - U'(y)] + i^b[\partial_\omega\gamma - \theta], \tag{149}$$

$$(\rho - \mathcal{A}_t)\gamma_t - \partial_t\gamma_t = \mathcal{U}_t + \mu_t y_t + \varphi_t(\omega - y_t) - \theta_t T_t, \qquad \psi_t \equiv 0, \tag{150}$$

$$\theta_t = \int \partial_\omega\gamma_t\,dG_t, \qquad \dot{\theta}_t = (\rho - r_t^b)\theta_t - \varphi_t, \qquad \mu_t - \varphi_t = i_t^b\theta_t, \tag{151}$$

$$\int(\omega - y_t)(\partial_\omega\gamma_t - \theta_t)dG_t = \int \eta_t\,u'(c_t)\,\mathbf{1}^{\text{int}}\,d\omega, \tag{152}$$

$$\int \omega\,\partial_\omega\gamma_t\,dG_t \;=\; \theta_t\,A_t \qquad (\delta g_t^M:\ \text{optimal-inflation condition}). \tag{153}$$

Proposition 9 (SPE discretion = commitment with $\psi \equiv 0$)**.** *Under subgame-perfect discretion the optimality system is* (149)–(153)*: the commitment system of Proposition 8 with* $\psi \equiv 0$ *(hence* $\Psi = \Psi^b = 0$

and $R_t = \int \omega \, \partial_\omega \gamma_t \, dG_t$*).*

C.5 Spatial boundary, terminal, and initial conditions

The adjoint equations (26)–(33) of Section 3.2 require boundary conditions at $\omega = \underline{b}$. For the backward HJB (26), the mass variation $\delta L/\delta m_s(z) = 0$ gives a backward HJB for γ_s at the constraint, with two simplifications relative to the interior: agents there have $a_s = 0$ and cannot make lumpy consumption ($y_s(\underline{b}, z) = 0$), so the Poisson jump term drops; and the effective discount rate is $\rho + \alpha$, since agents at the mass point cycle through Poisson meetings at rate α without changing their wealth. The social value at the constraint is governed by the flow welfare $\mathcal{U}_s(\underline{b}, z)$, the fiscal cost of transfers, the market-clearing cost $\varphi_s \underline{b}$ (with $y_s = 0$), and the productivity-transition dynamics, all discounted at $\rho + \alpha$. For the forward KFE (33), the condition $\delta L/\delta W_s(\underline{b}, z) = 0$ collects the ω-direction integration-by-parts residuals at $\omega = \underline{b}^+$: a contribution from the HJB wealth drift (the shadow price of the promise times the exit rate $\dot{\omega}_s(\underline{b}, z)$), shadow-cost terms from the flow and augmented-lumpy implementability constraints, and a contribution from the endogenous full-depletion threshold $\underline{\omega}_s(z)$, whose response to $\partial_\omega W_s(\underline{b}, z)$ alters the Poisson mass flow of agents depleting into the constraint.

The terminal and initial conditions of Proposition 1 follow from the time-boundary residuals. Varying G_{τ_1} and integrating the KFE Lagrange terms by parts in time generates the residual

$$e^{-\rho\tau_1}\left[\int \left(\gamma_{\tau_1}(x) - \frac{\delta \mathcal{W}_{\tau_1}}{\delta g_{\tau_1}(x)}\right) g_{\tau_1} \, dx + \sum_z \left(\gamma_{\tau_1}(\underline{b}, z) - \frac{\delta \mathcal{W}_{\tau_1}}{\delta m_{\tau_1}(z)}\right) m_{\tau_1}(z)\right] = 0,$$

which for arbitrary (g_{τ_1}, m_{τ_1}) gives the terminal condition on γ_{τ_1} (the functional derivative evaluated as $\delta \mathcal{W}_{\tau_1}/\delta g_{\tau_1}$ in the interior and $\delta \mathcal{W}_{\tau_1}/\delta m_{\tau_1}$ at the constraint). Varying W_0 and integrating the HJB multiplier term $-e^{-\rho s}\psi_s g_s \partial_s W_s$ by parts over $[0, \tau_1]$ gives the dual residual at $s = 0$, yielding the initial condition $\psi_0(x) = 0$.

In the limit $\tau_1 \to \infty$, these terminal and initial conditions are replaced by transversality conditions on the co-states:

$$\lim_{t\to\infty} e^{-\rho t}\, \theta_t \, A_t = 0, \qquad \lim_{t\to\infty} e^{-\rho t} \int \gamma_t(x) \, dG_t(x) = 0, \qquad \lim_{t\to\infty} e^{-\rho t} \int \psi_t(x) \, dx = 0, \tag{154}$$

with $A_t = \phi_t(M_t + B_t)$. The first is the planner's analog of the government no-Ponzi constraint; the second and third ensure the present-value contributions of the wealth distribution and past promises do not explode. Together with the household transversality conditions, they select the unique non-explosive solution to the optimality system.

D Proof of Proposition 5

This appendix proves Proposition 5, deriving the separated optimality system using the discretize-first approach taken in Appendix C. The discretization method—the discrete generator and its adjoints, the upwind derivative, and time summation-by-parts—is identical to the consolidated problem and is not repeated; this appendix only discusses what changes when the consolidated government budget splits into separate central-bank and Treasury budgets linked by a fiscal rule. Importantly, it shows the joint continuum limit recovers the conditions of Section 4.2 and those in Proposition 5.

D.1 The separated, detrended-nominal economy

D.1.1 Change of variables and the household block

Detrend by the money stock M_t, with $\tilde{\omega} = \omega/\phi_t$ as household nominal wealth,

$$\hat{\omega}_t \equiv \frac{\tilde{\omega}}{M_t} = \frac{\omega}{\phi_t M_t}, \quad Q_t \equiv \phi_t M_t, \quad \beta_t^T \equiv \frac{B_t^T}{M_t}, \quad \beta_t^{CB} \equiv \frac{B_t^{CB}}{M_t}, \quad \beta_t \equiv \beta_t^T - \beta_t^{CB}, \quad g_t^M \equiv \frac{\dot{M}_t}{M_t}, \tag{155}$$

so real wealth is $\omega = Q_t\hat{\omega}$, bonds in the hands of the public are $B_t = B_t^T - B_t^{CB}$ (so $\beta_t = B_t/M_t$), and the real returns are

$$r_t^a = \frac{\dot{\phi}_t}{\phi_t} = \frac{\dot{Q}_t}{Q_t} - g_t^M, \qquad r_t^b = r_t^a + i_t^b. \tag{156}$$

Relative to the consolidated case, we separate Treasury bonds outstanding, β^T, from central bank holdings, β^{CB}, relative to the money supply. The household drift, value function, and first-order conditions are identical to the consolidated case — the household block takes as input only the prices (r^a, i^b) and the real transfer T:

$$\dot{\hat{\omega}} = \hat{s}_t(\hat{\omega}, z) \equiv i_t^b\left(\hat{\omega} - \frac{y_t}{Q_t}\right) + \frac{z + T_t - c_t}{Q_t} - g_t^M\hat{\omega}, \qquad T_t = \bar{T} + \kappa(\bar{b} - Q_t\beta_t^T). \tag{157}$$

The one new feature relative to the consolidated drift is that the rule makes T_t a function of the aggregate state $Q_t\beta_t^T = \phi_t B_t^T$: the transfer is no longer a free planner control but a function of detrended Treasury debt. This single dependence is the source of every separated-specific term below. We maintain the fixed borrowing boundary $\hat{\omega} = 0$ ($\underline{b} = 0$).

D.1.2 Two budgets, market clearing, the fiscal rule

With $R_t = 0$, the detrended Treasury and central-bank flow budgets are

$$\text{Treasury } (\zeta): \quad Q_t(\dot{\beta^T}_t + g_t^M \beta_t^T) - T_t - i_t^b Q_t \beta_t^T = 0, \qquad T_t = \bar{T} + \kappa(\bar{b} - Q_t \beta_t^T), \tag{158}$$

$$\text{Central bank } (\vartheta): \quad Q_t(\dot{\beta^{CB}}_t + g_t^M \beta_t^{CB}) - i_t^b Q_t \beta_t^{CB} - Q_t g_t^M = 0, \tag{159}$$

where we used $\phi \dot{B} = Q(\dot{\beta} + g^M \beta)$ and $\phi \dot{M} = Q g^M$. Summing (158)–(159) (with $\beta = \beta^T - \beta^{CB}$) recovers the consolidated budget. The money- and bond-market clearing conditions are

$$\int y_t \, d\hat{F}_t = Q_t, \qquad \int (Q_t \hat{\omega} - y_t) \, d\hat{F}_t = Q_t \beta_t = Q_t(\beta_t^T - \beta_t^{CB}). \tag{160}$$

Total real public liabilities are $A_t = Q_t(1 + \beta_t)$, with $A^a = Q$ (money), $A^b = Q\beta$ (bonds in the hands of the public), $A^{b,T} = Q\beta^T = \phi B^T$ (Treasury debt), and $A^{b,CB} = Q\beta^{CB} = \phi B^{CB}$ (CB bond holdings); note $A^b = A^{b,T} - A^{b,CB}$ and $A = A^a + A^b = A^a - A^{b,CB} + A^{b,T}$.

D.1.3 The central bank's problem and Lagrangian

Controls, states, and objective. The planner controls, per slot, the household objects $(c, y, V, \hat{F})$ subject to implementability; the static price-like controls (Q_t, i_t^b), money growth g_t^M, its bond holdings β_t^{CB}, and the Treasury-debt state β_t^T, whose law of motion is the Treasury budget (158). There is no free transfer control. The state variables are the detrended aggregate objects $(\hat{F}_t, \beta_t^T, \beta_t^{CB})$. The objective follows similarly to the consolidated case.

$$\mathcal{W}_0(\hat{F}_0, \beta_0^T, \beta_0^{CB}) = \max \; \mathbb{E}_0 \left[\int_0^{\tau_1} e^{-\rho s} \int \mathcal{U}_s \, d\hat{F}_s \, ds \; + \; e^{-\rho \tau_1} \, \mathcal{W}_{\tau_1}(\hat{F}_{\tau_1}, \beta_{\tau_1}^T, \beta_{\tau_1}^{CB}) \right], \tag{161}$$

where $\mathcal{U}_s = u(c_s) + \alpha U(y_s)$ is the flow utility.

Lagrangian. Attaching the distributional multipliers $\hat{\psi}_s, \hat{\chi}_s, \hat{\eta}_s, \hat{\gamma}_s$ (HJB, c-Euler, y-Euler, KFE) and the scalars ζ_s (Treasury budget), ϑ_s (CB budget), μ_s (money clearing), φ_s (bond-market clear-

ing), the within-slot Lagrangian is

$$\begin{aligned}\mathcal{L} = \int_0^{\tau_1} e^{-\rho s}\Big\{ & \int \mathcal{U}_s\, d\hat{F}_s + \int \hat{\psi}_s[-\rho V_s + \mathcal{U}_s + \partial_s V_s + \hat{\mathcal{A}}_s V_s]d\hat{\omega}dz \\ & + \int \hat{\chi}_s[u'(c_s) - \frac{1}{Q_s}\partial_{\hat{\omega}} V_s]d\hat{\omega}dz \\ & + \int \hat{\eta}_s[\alpha U'(y_s) - \frac{1}{Q_s}(\alpha\partial_{\hat{\omega}} V_s(\hat{\omega} - y_s/Q_s, z) + i_s^b \partial_{\hat{\omega}} V_s)]\mathbf{1}_s^{\text{int}} d\hat{\omega}dz \\ & + \int \hat{\gamma}_s[\hat{\mathcal{A}}_s^* \hat{f}_s - \partial_s \hat{f}_s]d\hat{\omega}dz \\ & + \zeta_s[Q_s(\dot{\beta^T}_s + g_s^M \beta_s^T) - T_s - i_s^b Q_s \beta_s^T] + \vartheta_s[Q_s g_s^M + i_s^b Q_s \beta_s^{CB} - Q_s(\dot{\beta^{CB}}_s + g_s^M \beta_s^{CB})] \\ & + \mu_s[\int y_s\, d\hat{F}_s - Q_s] + \varphi_s[\int (Q_s\hat{\omega} - y_s)\, d\hat{F}_s - Q_s(\beta_s^T - \beta_s^{CB})]\Big\}\, ds \\ & + e^{-\rho\tau_1}\mathcal{W}_{\tau_1}(\hat{F}_{\tau_1}, \beta_{\tau_1}^T, \beta_{\tau_1}^{CB}). \end{aligned} \tag{162}$$

The same hat-scalings hold as in the consolidated case: $\hat{\psi} = Q\psi, \hat{\chi} = Q\chi, \hat{\eta} = Q\eta$, $\hat{\gamma}$ unscaled, and $\vartheta, \zeta, \mu, \varphi, i^b$ are real-unit scalars carried unchanged. Also, the detrended economy continues to yield a slot-locality property.

Lemma 11 (Slot locality)**.** *Every slot-n term of* (162) *depends only on slot-n objects* $(c_n, y_n, V_n, \hat{f}_n, Q_n, i_n^b, g_n^M, \beta_n^{CB}, \beta_{n+1}^{CB}, \beta_n^T, \beta_{n+1}^T)$ *and on V_{n+1}. The fiscal rule enters only through $Q_n\beta_n^T$, which is slot-local. No slot-n term references a successor price Q_m, ϕ_m for $m \geq n+1$.*

Proof. Identical to the consolidated problem: the drift (157), the budgets (158)–(159), clearing (160), and the rule all involve slot-n aggregates only, and $\hat{f}_n$ is nominal-detrended so needs no price to interpret. □

D.2 Discretized Lagrangian

The discretization is identical to the consolidated case: time into steps dt with $\partial_t x_t \to (x_{n+1} - x_n)/dt$ for any variable x_t; the wealth grid with weighted inner product $\langle h, g\rangle = \sum_i h_i g_i \Delta_i = h^\top D_\Delta g$ and $\int \Psi\, d\hat{F}_n \approx \langle \Psi, \hat{f}_n\rangle$; the upwind generator $\mathbf{A}_n = \mathbf{S}_n^{\text{drift}} + \alpha(\hat{\mathbf{J}}_n - \mathbf{I}) + \mathbf{A}_z$ with upwind derivative $\mathbf{D}_n^{\hat{\omega}}$ and adjoint $\mathbf{A}_n^* = D_\Delta^{-1}\mathbf{A}_n^\top D_\Delta$; and the time summation-by-parts identity.

Discretizing (162) gives the slot-n integrand $\mathcal{I}_n$ as the sum of the five household pieces (welfare, HJB, c-Euler, y-Euler, KFE), which are verbatim from the consolidated case: $\mathcal{I}_n^{\text{welfare}}, \ldots, \mathcal{I}_n^{\text{KFE}}$, and the four separated aggregate pieces

$$\mathcal{I}_n^{\text{TreasBC}} = \zeta_n\Big[Q_n g_n^M \beta_{n+1}^T + Q_n \frac{\beta_{n+1}^T - \beta_n^T}{dt} - T_n\langle \mathbf{1}^J, \hat{f}_n\rangle - i_n^b Q_n \beta_n^T\Big], \quad T_n = \bar{T} + \kappa(\bar{b} - Q_n\beta_n^T), \tag{163}$$

$$\mathcal{I}_n^{\text{CBbc}} = \vartheta_n\Big[Q_n g_n^M + i_n^b Q_n \beta_n^{CB} - Q_n g_n^M \beta_{n+1}^{CB} - Q_n \frac{\beta_{n+1}^{CB} - \beta_n^{CB}}{dt}\Big], \tag{164}$$

$$\mathcal{I}_n^{\text{moneyMC}} = \mu_n[\langle y_n, \hat{f}_n\rangle - Q_n], \qquad \mathcal{I}_n^{\text{bondMC}} = \varphi_n[\langle Q_n\hat{\boldsymbol{\omega}} - y_n, \hat{f}_n\rangle - Q_n(\beta_n^T - \beta_n^{CB})]. \tag{165}$$

The drift vector $\hat{s}_n$ (in the household HJB and KFE) now carries the fiscal rule, $T_n = \bar{T} + \kappa(\bar{b} - Q_n\beta_n^T)$, so $\hat{s}_n$ depends on the aggregate $Q_n\beta_n^T$ in addition to $(c_n, y_n, Q_n, i_n^b, g_n^M)$. All hat-scalings and the real-translation dictionary ($V = W(Q\hat{\omega})$, $\hat{f} = Qg$, $\langle \Psi, \hat{f} \rangle = \int \Psi \, dG$, $\mathbf{D}^{\hat{\omega}} V \leftrightarrow Q\partial_\omega W$, $\hat{\psi} = Q\psi$, etc.) are exactly as in the consolidated case.

D.3 Full commitment ($\lambda \to 0$)

As in the consolidated case, sum by parts both the KFE term ($\hat{\gamma}$) and the HJB term ($\hat{\psi}$) in time, with predecessor convention $\hat{\psi}_{-1} = 0$; the budgets' own $\beta^{\dot{T}}, \beta^{\dot{CB}}$ differences are handled directly in $\delta\beta_{n+1}^T, \delta\beta_{n+1}^{CB}$. We read the FOCs off the post-SBP integrand.

D.3.1 Household-block variations: δV_n, δc_n, δy_n, $\delta \hat{f}_{n+1}$

The four household variations are unchanged from the consolidated problem except for a single substitution in the γ-equation, because the household integrands are identical and the only government multiplier touching them is ϑ in the lumpy wedge (now on the CB's budget constraint). These are repeated here to completeness.

Private valuations, δV_n. V_n enters the post-SBP HJB constant pieces $\langle \hat{\psi}_n, -\rho V_n + \mathbf{A}_n V_n \rangle$, the time-SBP recursion, and the two Euler brackets. Collecting coefficients gives the $\hat{\psi}$-recursion

$$\frac{\hat{\psi}_n - e^{\rho dt}\hat{\psi}_{n-1}}{dt} = \mathbf{A}_n^* \hat{\psi}_n - \rho\hat{\psi}_n - \frac{1}{Q_n}(\mathbf{D}_n^{\hat{\omega}})^* \Big[\hat{\chi}_n + (\alpha \hat{\mathbf{J}}_n^* + i_n^b \mathbf{I})\mathrm{diag}(\mathbf{1}_n^{\mathrm{int}})\hat{\eta}_n\Big], \tag{166}$$

where $(\mathbf{D}_n^{\hat{\omega}})^*$ and $\hat{\mathbf{J}}_n^*$ are the D_Δ-adjoints of the upwind derivative and the jump matrix (Appendix C). Translating ($\hat{\psi} = Q\psi, \hat{\chi} = Q\chi, \hat{\eta} = Q\eta$) and taking $dt, \Delta \to 0$ gives the forward KFE

$$\partial_t \psi = \mathcal{A}^*[\psi] + \partial_\omega \chi + \partial_\omega\Big[\frac{\alpha\, \eta(\omega^+, z)}{1 - \partial_\omega y(\omega^+, z)}\Big] + \partial_\omega[i^b \eta] \quad (\to (33)). \tag{167}$$

Flow consumption, δc_n. The discretized FOC is

$$\hat{\chi}_n = \mathrm{diag}(u''(c_n))^{-1}\mathrm{diag}(\hat{f}_n)\Big[\frac{1}{Q_n}\mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n - u'(c_n)\Big], \tag{168}$$

with continuum limit

$$\chi = \frac{\partial_\omega \gamma - u'(c)}{u''(c)}\, g \quad (\to (31)). \tag{169}$$

Lumpy consumption, δy_n. Restricting to the interior ($\mathbf{1}_{n,j}^{\mathrm{int}} = 1$), the discretized FOC is

$$0 = \alpha \hat{f}_{n,j} U'(y_{n,j}) + \mathcal{H}_{n,j}\hat{\eta}_{n,j} - \frac{\hat{f}_{n,j}}{Q_n}[i_n^b(\mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n)_j + \alpha[\mathcal{D}_n^{\mathrm{post}}\hat{\gamma}_n]_j] + (\mu_n - \varphi_n)\hat{f}_{n,j}, \tag{170}$$

with curvature $\mathcal{H}_{n,j} = \alpha[U''(y_{n,j}) + \frac{1}{Q_n^2}[\mathcal{D}_n^{\text{post}}\mathbf{D}_n^{\hat{\omega}}V_n]_j]$ and $\mathcal{D}_n^{\text{post}}$ the post-jump interpolation-slope operator (Appendix C). Solving for the multiplier,

$$\hat{\eta}_{n,j} = \frac{\hat{f}_{n,j}\,\tilde{\eta}_{n,j}}{\mathcal{H}_{n,j}}, \qquad \tilde{\eta}_{n,j} = \alpha\Big[\frac{1}{Q_n}[\mathcal{D}_n^{\text{post}}\hat{\gamma}_n]_j - U'(y_{n,j})\Big] + \frac{i_n^b}{Q_n}(\mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n)_j + (\varphi_n - \mu_n). \tag{171}$$

Taking $dt, \Delta \to 0$ and substituting $\mu - \varphi = i^b\vartheta$ (from (193)) gives the lumpy consumption wedge

$$\tilde{\eta} = \alpha[\partial_\omega\gamma(\omega - y, z) - U'(y)] + i^b[\partial_\omega\gamma(\omega, z) - \vartheta] \quad (\to (32)), \tag{172}$$

with ϑ the central bank's shadow price. All three derivations above are identical to the consolidated problem (Appendix C): c, y, V enter only household terms and the two market-clearing conditions, so the budget split does not touch them (ϑ is merely reinterpreted as the CB's shadow price).

Distributional state, $\delta\hat{f}_{n+1}$. This is the only change from the consolidated case related to household objects. The Treasury transfer term $-\zeta_{n+1}T_{n+1}\langle\mathbf{1}^J, \hat{f}_{n+1}\rangle$ now holds the multiplier ζ associated with the Treasury budget (163). Setting $\partial\mathcal{L}^C/\partial\hat{f}_{n+1} = 0$ and dividing out $e^{-\rho t_{n+1}}dt\,D_\Delta$ gives the discretized FOC

$$0 = \mathcal{U}_{n+1} + \mathbf{A}_{n+1}\hat{\gamma}_{n+1} + \frac{\hat{\gamma}_{n+1} - e^{\rho dt}\hat{\gamma}_n}{dt} + \mu_{n+1}y_{n+1} + \varphi_{n+1}(Q_{n+1}\hat{\omega} - y_{n+1}) - \zeta_{n+1}T_{n+1}\mathbf{1}^J. \tag{173}$$

As $dt, \Delta \to 0$ the difference $(\hat{\gamma}_{n+1} - e^{\rho dt}\hat{\gamma}_n)/dt \to \partial_t\hat{\gamma} - \rho\hat{\gamma}$, and translating to real coordinates ($Q\hat{\omega} \leftrightarrow \omega$) gives the backward γ-HJB

$$\rho\gamma_s = \mathcal{U}_s - \zeta_s T_s + \mu_s y_s + \varphi_s(\omega - y_s) + \mathcal{A}_s[\gamma_s] + \partial_t\gamma_s \quad (\to (54)). \tag{174}$$

D.3.2 Aggregate variations: $\delta\beta_{n+1}^{CB}$, $\delta\beta_{n+1}^{T}$, δg_n^M, δi_n^b, δQ_n

Central-bank bond holdings, $\delta\beta_{n+1}^{CB}$: Varying bond holdings of the central bank, β_{n+1}^{CB} appears in the slot-n CB budget (164), in the slot-(n+1) CB budget, and in slot-(n+1) bond-market clearing. Collecting $\partial\mathcal{L}^C/\partial\beta_{n+1}^{CB} = 0$ with the discount factors,

$$\vartheta_n Q_n(1 + g_n^M dt) = e^{-\rho dt}Q_{n+1}[\vartheta_{n+1}(1 + i_{n+1}^b dt) + \varphi_{n+1}dt]. \tag{175}$$

Notice, this takes the same form as the consolidated variation with respect to bonds, $\delta\beta$, in (126). Expanding to $O(dt)$, using $r^a = \dot{Q}/Q - g^M$, and taking limits yields

$$\dot{\vartheta}_t = (\rho - r_t^a - i_t^b)\vartheta_t - \varphi_t = (\rho - r_t^b)\vartheta_t - \varphi_t \quad (\to (56)). \tag{176}$$

Treasury debt, $\delta\beta^T_{n+1}$: β^T_{n+1} enters four channels, each carrying the discount factor of the slot it sits in:

(i) The slot-n Treasury budget (163): the $g^M\beta^T_{n+1}$ and time-difference terms give coefficient $\zeta_n(Q_n g^M_n + Q_n/dt)$.

(ii) The slot-(n+1) Treasury budget: entering the time difference ($-Q_{n+1}/dt$), the interest term $-i^b_{n+1}Q_{n+1}\beta^T_{n+1}$ ($-i^b_{n+1}Q_{n+1}$), and the transfer $-T_{n+1}$ with $T_{n+1} = \bar{T} + \kappa(\bar{b} - Q_{n+1}\beta^T_{n+1})$, so $\partial(-T_{n+1})/\partial\beta^T_{n+1} = +\kappa Q_{n+1}$; total coefficient $\zeta_{n+1}(-Q_{n+1}/dt - i^b_{n+1}Q_{n+1} + \kappa Q_{n+1})$.

(iii) The slot-(n+1) bond-market clearing: $-\varphi_{n+1}Q_{n+1}$.

(iv) The slot-(n+1) household drift: T_{n+1} enters $\hat{s}_{n+1}$ through $(z + T - c)/Q$, so $\partial\hat{s}_{n+1}/\partial T_{n+1} = 1/Q_{n+1}$ and, with $\partial T_{n+1}/\partial\beta^T_{n+1} = -\kappa Q_{n+1}$, the chain rule gives $\partial\hat{s}_{n+1}/\partial\beta^T_{n+1} = -\kappa$ (a scalar, uniform in x). Through the $\hat{\psi}$-HJB term $\langle\hat{\psi}, \mathbf{A}V\rangle$ and the post-SBP KFE term $\langle\mathbf{A}\hat{\gamma}, \hat{f}\rangle$ this contributes coefficient $-\kappa\,\Pi_{n+1}$, where $\Pi_n \equiv \langle\hat{\psi}_n, \mathbf{D}^{\hat{\omega}}_n V_n\rangle + \langle\mathbf{D}^{\hat{\omega}}_n\hat{\gamma}_n, \hat{f}_n\rangle$ is the household-drift inner product.

Setting $\partial\mathcal{L}^C/\partial\beta^T_{n+1} = 0$:

$$0 = e^{-\rho t_n}dt\,\zeta_n\left(Q_n g^M_n + \frac{Q_n}{dt}\right) + e^{-\rho t_{n+1}}dt\left[\zeta_{n+1}\left(-\frac{Q_{n+1}}{dt} - i^b_{n+1}Q_{n+1} + \kappa Q_{n+1}\right) - \varphi_{n+1}Q_{n+1} - \kappa\,\Pi_{n+1}\right]. \tag{177}$$

Reduce (177) to a co-state recursion. Divide by $e^{-\rho t_n}$, write $t_{n+1} = t_n + dt$, multiply by dt, and distribute Q_{n+1}:

$$\zeta_n Q_n(1 + g^M_n dt) = e^{-\rho dt}[\zeta_{n+1}Q_{n+1}(1 + (i^b_{n+1} - \kappa)dt) + \varphi_{n+1}Q_{n+1}dt + \kappa\Pi_{n+1}dt]. \tag{178}$$

This form mimics the central-bank recursion (175) with ζ instead of ϑ, the extra $-\kappa$ from the rule's effect on the slot-(n+1) budget (channel ii), and the discrete household-drift source $\kappa\Pi_{n+1}$ (channel iv).

Now take limits. The finite differences expand to $O(dt)$ as in (175)→(176) above. The household-drift inner product Π_n has an explicit continuum limit. Writing $\langle h, g\rangle = \sum_i h_i g_i \Delta_i$ and using $\hat{\psi} = Q\psi$, $\mathbf{D}^{\hat{\omega}}V \to Q\partial_\omega W$, $\mathbf{D}^{\hat{\omega}}\hat{\gamma} \to Q\partial_\omega\gamma$, $\hat{f} = Qg$, and the cell-volume rule $\Delta_i \to d\hat{\omega}\,dz = d\omega\,dz/Q$, each inner product picks up exactly one factor of Q:

$$\Pi_n \xrightarrow{dt,\Delta\to 0} Q_n\int(\psi_n\,\partial_\omega W_n + \partial_\omega\gamma_n\,g_n)\,dx, \tag{179}$$

so the channel-(iv) coefficient $-\kappa\,\Pi_n \to -\kappa Q_n\,\theta_n$, where

$$\theta_s \equiv \int(\psi_s\,\partial_\omega W_s + \partial_\omega\gamma_s\,g_s)\,dx = \int\partial_\omega\gamma_s\,dG_s + \int\psi_s\,\partial_\omega W_s\,d\omega. \tag{180}$$

The scalar θ_s is the aggregate social marginal value of wealth—the social value of a uniform lump-sum transfer, unweighted by ω (the wealth-weighted object is R). Since the separated planner has no free-transfer FOC, θ_s equals neither ϑ_s nor ζ_s (in the consolidated problem the transfer FOC would set $\vartheta = \theta$). Collectively, the limit gives

$$\dot{\zeta}_t = (\rho - r_t^b)\zeta_t - \varphi_t + \kappa(\zeta_t - \theta_t) \quad (\to (57)). \tag{181}$$

The fiscal-rule feedback $\kappa(\zeta_t - \theta_t)$ has two pieces: $+\kappa\zeta$ from the rule's effect on the Treasury budget, $-\kappa\theta$ from its effect on households. At $\kappa = 0$ it vanishes and ζ obeys the same ODE as ϑ in (176).

Money growth δg_n^M: Money growth g_n^M enters: the drift with $\partial \hat{s}_n / \partial g_n^M = -\hat{\boldsymbol{\omega}}$ (dilution), weighted by the $\hat{\psi}$-HJB and the post-SBP KFE; the Treasury budget (163) with coefficient $\zeta_n Q_n \beta_{n+1}^T$; and the CB budget (164) with coefficient $\vartheta_n[Q_n - Q_n\beta_{n+1}^{CB}] = \vartheta_n Q_n(1 - \beta_{n+1}^{CB})$. The FOC is

$$\langle \mathrm{diag}(\hat{\boldsymbol{\omega}})\mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n, \hat{f}_n\rangle + \langle \hat{\psi}_n, \mathrm{diag}(\hat{\boldsymbol{\omega}})\mathbf{D}_n^{\hat{\omega}}V_n\rangle = \vartheta_n Q_n(1 - \beta_{n+1}^{CB}) + \zeta_n Q_n \beta_{n+1}^T. \tag{182}$$

Taking limits, ($\langle \mathrm{diag}(\hat{\boldsymbol{\omega}})\mathbf{D}\hat{\gamma}, \hat{f}\rangle \to \int \omega\, \partial_\omega \gamma \, dG$ and likewise for the term involving $\hat{\psi}$; $Q \to A^a$, $Q\beta^{CB} \to A^{b,CB}$, $Q\beta^T \to A^{b,T}$),

$$R_t \equiv \int \omega\, \partial_\omega \gamma_t \, dG_t + \int \omega\, \psi_t\, \partial_\omega W_t \, d\omega \;=\; \vartheta_t(A_t^a - A_t^{b,CB}) + \zeta_t A_t^{b,T} \;\equiv\; \Lambda_t \quad \forall t, \tag{183}$$

equivalent to (59). This is the analog of the consolidated optimality condition for money growth. A uniform money-growth dilution lowers every household's real claims (welfare effect R_t); the revenue Λ_t accrues along the two budgets — the central bank gains on net money issuance $A^a - A^{b,CB}$, valued at ϑ (seigniorage net of the cost of remunerating its own bond holdings), and the Treasury gains on its diluted debt $A^{b,T}$, valued at ζ.

Nominal bond rate δi_n^b: The bond rate, i_n^b, enters: the drift with $\partial \hat{s}_n / \partial i_n^b = \hat{\boldsymbol{\omega}} - y_n/Q_n$ (HJB and KFE); the y-Euler bracket with $-\frac{1}{Q_n}\mathbf{D}_n^{\hat{\omega}}V_n$; the Treasury budget (163) with coefficient $-\zeta_n Q_n \beta_n^T$ (the Treasury pays i^b on all of B^T); and the CB budget (164) with coefficient $+\vartheta_n Q_n \beta_n^{CB}$ (the CB earns i^b on B^{CB}). Setting $\partial \mathcal{L}^C / \partial i_n^b = 0$:

$$\langle \mathrm{diag}(\hat{\boldsymbol{\omega}} - y_n/Q_n)\mathbf{D}_n^{\hat{\omega}}\hat{\gamma}_n, \hat{f}_n\rangle + \langle \hat{\psi}_n, \mathrm{diag}(\hat{\boldsymbol{\omega}} - y_n/Q_n)\mathbf{D}_n^{\hat{\omega}}V_n\rangle - \frac{1}{Q_n}\langle \hat{\eta}_n, \mathbf{1}_n^{\mathrm{int}}\mathbf{D}_n^{\hat{\omega}}V_n\rangle + \vartheta_n Q_n \beta_n^{CB} - \zeta_n Q_n \beta_n^T = 0. \tag{184}$$

In the limit ($b = \omega - y$, $Q\beta^{CB} \to A^{b,CB} = \phi B^{CB}$, $Q\beta^T \to A^{b,T} = \phi B^T$),

$$\int b_s(\psi_s\, \partial_\omega W_s + \partial_\omega \gamma_s\, g_s)\, dx - \int \eta_s\, u'(c_s)\, dx \;=\; \zeta A^{b,T} - \vartheta A^{b,CB}, \tag{185}$$

equivalent to (58). The right-hand side, $\zeta A^{b,T} - \vartheta A^{b,CB}$, is the Treasury's interest cost (valued at ζ) net of the CB's interest income (valued at ϑ). This is the net fiscal cost of raising i^b, and it is balanced against the household-side benefit on the left.

Real balances δQ_n: As in the consolidated problem, Q_n enters more channels than any other control. The drift derivative now carries a new rule term: with $T_n = \bar{T} + \kappa(\bar{b} - Q_n \beta_n^T)$,

$$\frac{\partial \hat{s}_{n,i}}{\partial Q_n} = \frac{i_n^b y_{n,i} - (z_i + T_n - c_{n,i})}{Q_n^2} \underbrace{- \frac{\kappa \beta_n^T}{Q_n}}_{\text{transfer-rule}}, \tag{186}$$

where the second term captures $\partial T_n / \partial Q_n = -\kappa \beta_n^T$, divided by Q_n. The jump location, the c- and y-Euler $1/Q$ factors, and both clearing conditions enter exactly as in the consolidated problem; the budgets enter through (163)–(164).

Setting $\partial \mathcal{L}^C / \partial Q_n = 0$ and dividing out $e^{-\rho t_n} dt$ gives the raw discretized FOC

$$\begin{aligned}
0 = {} & \langle \mathrm{diag}(\partial_{Q_n} \hat{s}_n)\, \mathbf{D}_n^{\hat{\omega}} \hat{\gamma}_n, \hat{f}_n \rangle + \langle \hat{\psi}_n, \mathrm{diag}(\partial_{Q_n} \hat{s}_n)\, \mathbf{D}_n^{\hat{\omega}} V_n \rangle \\
& + \alpha \langle \mathrm{diag}(y_n / Q_n^2)\, \mathcal{D}_n^{\mathrm{post}} \hat{\gamma}_n, \hat{f}_n \rangle + \alpha \langle \hat{\psi}_n, \mathrm{diag}(y_n / Q_n^2)\, \mathcal{D}_n^{\mathrm{post}} V_n \rangle \\
& + \frac{1}{Q_n^2} \langle \hat{\chi}_n, \mathbf{D}_n^{\hat{\omega}} V_n \rangle + \langle \hat{\eta}_n, \mathbf{1}_n^{\mathrm{int}}\, \partial_{Q_n} E_n^y \rangle + \varphi_n [\langle \hat{\boldsymbol{\omega}}, \hat{f}_n \rangle - \beta_n] - \mu_n \\
& + \zeta_n \left[g_n^M \beta_{n+1}^T + \frac{\beta_{n+1}^T - \beta_n^T}{dt} + \kappa \beta_n^T \langle \mathbf{1}^J, \hat{f}_n \rangle - i_n^b \beta_n^T \right] \\
& + \vartheta_n \left[g_n^M + i_n^b \beta_n^{CB} - g_n^M \beta_{n+1}^{CB} - \frac{\beta_{n+1}^{CB} - \beta_n^{CB}}{dt} \right],
\end{aligned} \tag{187}$$

where $\beta_n = \beta_n^T - \beta_n^{CB}$, $\partial_{Q_n} \hat{s}_n$ is (186), and $\partial_{Q_n} E_{n,i}^y = \frac{1}{Q_n^2} [(\alpha \hat{\mathbf{J}}_n + i_n^b \mathbf{I}) \mathbf{D}_n^{\hat{\omega}} V_n]_i - \frac{\alpha y_{n,i}}{Q_n^3} [\mathcal{D}_n^{\mathrm{post}} \mathbf{D}_n^{\hat{\omega}} V_n]_i$ is the Q_n-derivative of the y-Euler bracket (as in Appendix C). The first three lines (drift, jump, c- and y-Euler, and clearing) coincide with the consolidated case; the last two lines replace the single consolidated budget with the Treasury (ζ_n) and central-bank (ϑ_n) budget components.

To consider the continuum limit, follow similar steps to the consolidated case. Multiply (187) by Q_n and translate to real variables. The household blocks (drift consolidated-piece, jump, χ, η, and the ψ/γ channels) reduce identically to the consolidated case via the pointwise commutation lemma $\omega \mathcal{A}[f] = \mathcal{A}[\omega f] + \alpha y f(\omega - y) - s^\omega f$ and the flow-derivative identity $\partial_\omega \mathcal{U} = (r^b - \partial_\omega s^\omega) \partial_\omega W + \alpha \partial_{\omega y} \partial_\omega W (\omega - y)$ (Lemmas 9–10): every flow, post-jump, χ, and η term cancels pairwise, the two time derivatives collect into $\dot{R} = \partial_t (\langle \omega \rangle_\gamma + \langle \omega \rangle_\psi)$, and the clearing line assembles (with $A = A^a + A^b$, $A^b = A^{b,T} - A^{b,CB}$) into $-\mu A^a - \varphi A^b$. Two pieces are new in the separated problem:

(i) the transfer-rule term $-\kappa \beta_n^T / Q_n$ in $\partial_{Q_n} \hat{s}_n$ (186), whose drift contribution is $-\kappa \beta_n^T [\langle \mathbf{D} \hat{\gamma}_n, \hat{f}_n \rangle + \langle \hat{\psi}_n, \mathbf{D} V_n \rangle] \to -\kappa\, A^{b,T}\, \theta$ (using $Q \beta^T = A^{b,T}$ and (180));

(ii) the two budget lines: after $\times Q_n$ and using each budget constraint to substitute its own bracket, the CB bracket is identically zero (the ϑ-channel contributes nothing), while the Treasury bracket equals T_n plus the rule's Q-derivative term $\kappa Q_n \beta_n^T$, so the ζ-channel contributes $\zeta_n(T_n + \kappa A^{b,T})$.

Collecting terms yields a similar co-state ODE to (143),

$$\dot{R} = \rho R + \Omega^{\text{sep}}, \qquad \Omega^{\text{sep}} = \zeta T - \mu A^a - \varphi(A^{b,T} - A^{b,CB}) + \underbrace{\zeta \kappa A^{b,T}}_{\text{rule in budget}} - \underbrace{\kappa A^{b,T} \theta}_{\text{transfer-rule drift}}. \tag{188}$$

The source term differs from the consolidated case (where there, $\Omega = \vartheta T - \mu A^a - \varphi A^b$). The term includes an additional component $\kappa A^{b,T}(\zeta - \theta)$ that captures how a marginal price-of-money change feeds through the fiscal rule: by revaluing real Treasury debt $A^{b,T} = \phi B^T$, it shifts the rule-determined transfer to every household, an effect priced at the Treasury's shadow value ζ on the budget side and at the social marginal value of wealth θ on the household side.

Collapse to $\mu - \varphi = i^b \vartheta$. As in the consolidated case, substitute the budgets and co-states into (188). Using the real-debt laws of motion $\dot{A}^a = (r^a + g^M)A^a$, $\dot{A^{b,T}} = r^b A^{b,T} + T$ (Treasury budget), $\dot{A^{b,CB}} = r^b A^{b,CB} + g^M A^a$ (CB budget), and the co-states (176) and (181), differentiate $\Lambda = \vartheta(A^a - A^{b,CB}) + \zeta A^{b,T}$ with respect to time:

$$\begin{aligned} \dot{\Lambda} = {} & [(\rho - r^b)\vartheta - \varphi](A^a - A^{b,CB}) + \vartheta(r^a A^a - r^b A^{b,CB}) \\ & + [(\rho - r^a + \kappa - i^b)\zeta - \varphi - \kappa\theta]A^{b,T} + \zeta(r^b A^{b,T} + T). \end{aligned} \tag{189}$$

Collecting (using $r^a - r^b = -i^b$ and $-r^a + \kappa - i^b + r^b = \kappa$),

$$\dot{\Lambda} = \vartheta[\rho(A^a - A^{b,CB}) - i^b A^a] - \varphi(A^a - A^{b,CB}) + \zeta[(\rho + \kappa)A^{b,T} + T] - \kappa\theta A^{b,T} - \varphi A^{b,T}. \tag{190}$$

Now form $\dot{\Lambda} - \rho\Lambda - \Omega^{\text{sep}}$ with $\rho\Lambda = \rho\vartheta(A^a - A^{b,CB}) + \rho\zeta A^{b,T}$ and Ω^{sep} from (188). The ζ-terms ($\rho A^{b,T}$, $\kappa A^{b,T}$, T), the $\kappa\theta A^{b,T}$ terms, and the $\rho\vartheta(A^a - A^{b,CB})$ terms cancel exactly; what remains is

$$\dot{\Lambda} - \rho\Lambda - \Omega^{\text{sep}} = -\vartheta i^b A^a - \varphi A^a + \mu A^a = A^a(\mu - \varphi - i^b\vartheta). \tag{191}$$

Since the δQ FOC is $\dot{R} = \rho R + \Omega^{\text{sep}}$, subtracting gives

$$\frac{d}{dt}(R - \Lambda) = \rho(R - \Lambda) - A^a(\mu - \varphi - i^b\vartheta). \tag{192}$$

The coefficient $-A^a \neq 0$, so δQ is not implied by δg^M. Imposing the static revaluation (183) ($R = \Lambda$ for all t, hence $\dot{R} = \dot{\Lambda}$) leaves

$$\mu_t - \varphi_t = i_t^b \, \vartheta_t \tag{193}$$

exactly as in the consolidated case. Combining (193) with the CB bond co-state (176) recovers the money co-state $\dot{\vartheta} = (\rho - r^a)\vartheta - \mu$.

Initial conditions. The initial conditions follow exactly to the consolidated case, with the exception of the state, here the triple $(\hat{F}_0, \beta_0^T, \beta_0^{CB})$.

D.3.3 Summary: the commitment optimality system

Collecting (167)–(172), (174), (176), (181), (183), (185), and (193), the separated commitment optimality system is collected in Table 5.

variation	discretized FOC	continuum limit	paper counterpart
δc	(168)	(169)	(31)
δy	(170)	(172)	(32)
δV	(166)	(167)	(33)
$\delta \hat{F}$	(173)	(174)	(54)
$\delta \beta^{CB}$	(175)	(176)	(56)
$\delta \beta^T$	(178)	(181)	(57)
δg^M	(182)	(183)	(59)
δi^b	(184)	(185)	(58)
δQ	(187)	(188)	(36)

Table 5: Cross-check of the separated commitment optimality system: each variation, its discretized FOC, the continuum limit, and the paper counterpart.

D.4 Discretion ($\lambda \to \infty$)

As in the consolidated case, the slot-s discretion planner sums by parts only the KFE, leaving the HJB raw. The δV_s variation then has the bare $-1/dt$ time term dominate, $[\mathbf{A}_s^* - \rho - 1/dt]\psi_s = \mathbf{b}_s$ with $\mathbf{b}_s = O(1)$, so $\psi_s = -dt\,\mathbf{b}_s + O(dt^2) \to 0$. By Lemma 11 every other variation is slot-local, so with $\psi \equiv 0$ each reproduces its commitment counterpart.

The only variations needing care are $\delta\beta_{s+1}^T$ and $\delta\beta_{s+1}^{CB}$. Under commitment the budget co-states ζ and ϑ are carried backward across the slot boundary; under discretion the slot-s planner instead sees these states only through the continuation value $\mathcal{W}_{s+1}$, so the two variations pick up the marginal continuation values $\partial\mathcal{W}_{s+1}/\partial\beta_{s+1}^T$ and $\partial\mathcal{W}_{s+1}/\partial\beta_{s+1}^{CB}$. The key point is that in a Markov-perfect equilibrium $\mathcal{W}_{s+1}$ is the successor's own optimized value, so the envelope

theorem applies: the derivative of an optimized value with respect to an inherited state equals the successor's co-state on that state's law of motion.

The law of motion of β^T is the Treasury budget (158) and that of β^{CB} is the central-bank budget (159), so the two derivatives are the successor's budget co-states ζ_{s+1} and ϑ_{s+1}, evaluated at its optimum—the same multipliers the commitment problem propagates backward across the slot boundary. The limit recursions (181) and (176) are therefore unchanged, as in the consolidated case with discretion. With $\psi \equiv 0$ the promise term drops everywhere it appeared, leaving $\theta_s = \int \partial_\omega \gamma_s \, dG_s$ and $R_s = \int \omega \, \partial_\omega \gamma_s \, dG_s$. The discretion optimality system is the continuum-limit conditions at $\psi \equiv 0$ (with their discretized counterparts in Table 5):

$$\chi = \frac{\partial_\omega \gamma - u'(c)}{u''(c)} \, g, \qquad \tilde{\eta} = \alpha[\partial_\omega \gamma(\omega - y, z) - U'(y)] + i^b[\partial_\omega \gamma - \vartheta], \tag{194}$$

$$\rho \gamma_s = \mathcal{U}_s - \zeta_s T_s + \mu_s y_s + \varphi_s(\omega - y_s) + \mathcal{A}_s[\gamma_s] + \partial_t \gamma_s, \qquad \psi_s \equiv 0, \tag{195}$$

$$\dot{\vartheta} = (\rho - r^b)\vartheta - \varphi, \qquad \dot{\zeta} = (\rho - r^b)\zeta - \varphi + \kappa(\zeta - \theta), \qquad \mu - \varphi = i^b \vartheta, \tag{196}$$

$$\int \omega \, \partial_\omega \gamma \, dG = \vartheta(A^a - A^{b,CB}) + \zeta A^{b,T}, \tag{197}$$

$$\int (\omega - y) \, \partial_\omega \gamma \, dG - \int \eta \, u'(c) \, dx = (\zeta B^T - \vartheta B^{CB})\phi, \tag{198}$$

with $\theta = \int \partial_\omega \gamma \, dG$.

E Numerical Algorithm

This appendix documents the computational procedure for the separated problem of Section 4 under discretion. The approach uses the discretization of the detrended model, outlined in Appendix D). The stationary optimum is computed by a nested bracketing scheme (Section E.2). Transition paths after an unanticipated aggregate shock are computed by stacking the discrete system over a truncated horizon and solving it with a Newton method whose Jacobian is assembled using the sequence-space method of Auclert et al. (2021) (Sections E.3–E.4). Throughout, ϑ_t is the multiplier on the central bank's budget, ζ_t the multiplier on the Treasury's, and $\theta_t \equiv \int \partial_\omega \gamma_t \, dG_t$ the social value of a uniform transfer; under discretion $\psi_t \equiv 0$ everywhere.

E.1 Discretization

State space. The idiosyncratic state (ω, z) is discretized with N_ω points for real wealth and N_z income states, $N = N_\omega N_z$ points in total. The baseline uses $N_\omega = 800$ on a non-uniform grid concentrated near the borrowing constraint (power-law spacing with exponent 3), an upper bound near thirty times annual output, and the two-state income process of Section 2.5.

Detrended-nominal implementation. All computations are carried out in detrended-nominal units (see Appendix D): household wealth is measured relative to the money stock. Real objects are recovered by scaling with real balances A_t^a. This removes the price-level path as a numerical unknown; only the time-0 revaluation survives as a scalar (Section E.3).

Household policies. At each grid point $i = (i_\omega, i_z)$ and date t, given the value function $\mathbf{W}_t$ and prices (r_t^a, i_t^b): flow consumption inverts the Euler equation, $c_{i,t} = (u')^{-1}(\partial_\omega W_{i,t})$, with forward and backward difference candidates resolved by upwinding. The lumpy policy is piecewise, $y_{i,t} = 0$ at the boundary mass point, $y_{i,t} = \omega_i - \underline{b}$ in the full-depletion region $\omega_i \leq \underline{\omega}_t(z)$, and otherwise the interior solution of the augmented Euler equation (13),

$$\alpha[U'(y) - \partial_\omega W_t(\omega_i - y, z)] = i_t^b \, \partial_\omega W_t(\omega_i, z), \tag{199}$$

with the full-depletion threshold $\underline{\omega}_t(z)$ solved per income state by a one-dimensional root-find on (20), and between-grid-point values of $\partial_\omega W_t$ obtained by linear interpolation in ω.

Generator and time stepping. The wealth drift $\dot{\omega}_{i,t} = r_t^b \omega_i - i_t^b y_{i,t} + z_i - c_{i,t} + T_t$ is discretized by the upwind scheme of Achdou et al. (2022), producing a sparse generator $\mathbf{A}_t \in \mathbb{R}^{N \times N}$ that collects the drift, the income transitions, and the Poisson lumpy jumps $\omega \to \omega - y_{i,t}$ (linear interpolation weights onto the grid); rows sum to zero. The value function steps backward by

implicit Euler,

$$\Big[(\rho + \frac{1}{\Delta t})\mathbf{I} - \mathbf{A}_t\Big]\mathbf{W}_t = \mathbf{u}_t + \frac{1}{\Delta t}\,\mathbf{W}_{t+\Delta t}, \tag{200}$$

an M-matrix system guaranteeing monotonicity; the distribution steps forward by $\mathbf{g}_{t+\Delta t} = (\mathbf{I} + \Delta t\,\mathbf{A}_t^\top)\,\mathbf{g}_t$ with unit-mass normalization. The stationary versions iterate (200) to a fixed point and solve $\mathbf{A}^\top\mathbf{g} = 0$. The backward equation for the social value γ_t below uses the same left-hand operator as (200), so one sparse factorization per date serves both solves.

Derivative moments. The planner-side objects $\partial_\omega\gamma_{i,t}$, $\partial_\omega c_{i,t}$, $\partial_\omega y_{i,t}$, and $\partial_{\omega\omega}W_{i,t}$ are computed by centered differences on the non-uniform grid (one-sided at the endpoints), with off-grid evaluation at post-depletion wealth $\omega_i - y_{i,t}$ by linear interpolation.

Implementability wedges. At each grid point, matching (31)–(32) with the central bank's shadow price in the portfolio term,

$$\tilde{\chi}_{i,t} = \partial_\omega\gamma_{i,t} - u'(c_{i,t}), \qquad \tilde{\eta}_{i,t} = \alpha[\partial_\omega\gamma_t(\omega_i - y_{i,t}, z) - U'(y_{i,t})] + i_t^b[\partial_\omega\gamma_{i,t} - \vartheta_t], \tag{201}$$

and the multiplier $\eta_{i,t} = \tilde{\eta}_{i,t}\,g_{i,t}/(\alpha[U''(y_{i,t}) + \partial_{\omega\omega}W_t(\omega_i - y_{i,t}, z)])$, defined on the interior region $\omega_i > \underline{\omega}_t(z)$ only.

E.2 Computing the stationary, separated discretionary optimum

Budget identities. The central bank chooses its remittance flow subject to the limit $0 \le R_t \le \bar{R}$, where the baseline assumes $\bar{R} = 0$. The marginal value of remitting a unit to the Treasury is $\zeta_t - \vartheta_t$. An interior choice would equate the two shadow prices and re-consolidate the budgets; by the steady state of (57), ζ_{ss} is a convex combination of ϑ_{ss} and θ_{ss}, so the gap $\zeta_{ss} - \vartheta_{ss}$ shares the sign of $\theta_{ss} - \vartheta_{ss}$, which is positive at every computed optimum, and the constrained choice sits at the limit, $R_{ss} = \bar{R}$, with Kuhn–Tucker multiplier $\zeta_{ss} - \vartheta_{ss} \ge 0$. The baseline sets $\bar{R} = 0$, so $R_{ss} = 0$. With $R_{ss} = \bar{R}$, the stationary fiscal rule and the two budgets reduce to

$$A_{ss}^{b,T} = \frac{\bar{T} + \kappa\bar{b} - \bar{R}}{\kappa - r_{ss}^b}, \qquad T_{ss} = \bar{R} - r_{ss}^b A_{ss}^{b,T}, \qquad A_{ss}^{b,CB} = \frac{r_{ss}^a A_{ss}^a + \bar{R}}{r_{ss}^b}. \tag{202}$$

Inner loop: bond-market clearing. Fix a candidate real money return r^a. For each candidate real bond return r^b on the admissible set $r^b \in [r^a, \rho)$, with $i^b = r^b - r^a \ge 0$, solve the stationary household block at (r^a, i^b, T) with T from (202), obtaining real balances A^a (money-market clearing), total wealth A, and household bond demand $A^{b,hh} = A - A^a$; then evaluate the clearing

residual

$$F(r^b; r^a) \equiv A^{b,hh} - (A^{b,T} - A^{b,CB}), \tag{203}$$

with $A^{b,T}$ and $A^{b,CB}$ from (202). Bisection solves the market-clearing $r^b(r^a)$.[43]

Co-states. At the clearing point, the stationary social value solves the backward equation (54) with sources $\mu\, y + \varphi\,(\omega - y)$ and a uniform term $-\zeta T$, where $\mu = (\rho - r^a)\vartheta$ and $\varphi = (\rho - r^b)\vartheta$. The distributional sources $\mu\, y + \varphi\,(\omega - y)$ are proportional to ϑ and vary across households. The remittance source $-\zeta T$ is uniform across households, so shifts γ by a constant. Since only gradient moments enter the optimality conditions, this term never needs to be tracked. Define $\hat{\gamma}^y$ and $\hat{\gamma}^b$ as the solutions of

$$(\rho\, \mathbf{I} - \mathbf{A}_{ss})\, \hat{\gamma}^y = \mathbf{y}_{ss}, \qquad (\rho\, \mathbf{I} - \mathbf{A}_{ss})\, \hat{\gamma}^b = \boldsymbol{\omega} - \mathbf{y}_{ss}, \tag{204}$$

where $\mathbf{A}_{ss}$ is the stationary generator of the household state. These represent present-value ledgers of money and bond holdings of a particular household. The term $\hat{\gamma}^y$ is the expected discounted stream of the household's future money holdings along its equilibrium path; the term $\hat{\gamma}^b$ is the stream of its future bond holdings. Linearity then gives

$$\gamma = \mathbf{W} + \mu\, \hat{\gamma}^y + \varphi\, \hat{\gamma}^b, \tag{205}$$

up to a uniform constant from the $-\zeta T$ term.

The central bank's shadow price. Solve for ϑ_{ss} using the stationary optimal-inflation condition (59),

$$\sum_i \omega_i\, \partial_\omega \gamma_{i,ss}\, g_{i,ss}\, \Delta_i \;=\; \vartheta_{ss} A^a_{ss} + \zeta_{ss} A^{b,T}_{ss} - \vartheta_{ss} A^{b,CB}_{ss}, \tag{206}$$

together with ζ_{ss} from the steady state of (57) and $\theta_{ss} = \sum_i \partial_\omega \gamma_{i,ss}\, g_{i,ss}\, \Delta_i$. By (205) both sides are affine in ϑ_{ss}, so (206) is a scalar linear equation solved exactly for ϑ_{ss} at each candidate r^a; ζ_{ss} and θ_{ss} then follow immediately.

[43] When evaluating, it is necessary to consider the pole around $r^b = 0$, if admissible. If $r^a < 0$, the position $A^{b,CB}_{ss}$ in (202) diverges as $r^b \to 0$. If $r^b < 0$ (when $r^a < 0$), then the feasible branch carries configurations in which the central bank is a large net creditor financed by seigniorage—the net-creditor case discussed after Proposition 6—and must be scanned separately.

Outer loop: the premium margin. With $(\vartheta_{ss}, \zeta_{ss}, \theta_{ss})$ set, the remaining condition is the bond-rate optimality condition (58), evaluated using

$$S(r^a) \equiv \sum_i (\omega_i - y_{i,ss})\, \partial_\omega \gamma_{i,ss}\, g_{i,ss}\, \Delta_i \;-\; \Big[\zeta_{ss} A^{b,T}_{ss} - \vartheta_{ss} A^{b,CB}_{ss} + \sum_{i:\,\omega_i > \underline{\omega}_{ss}(z)} \eta_{i,ss}\, u'(c_{i,ss})\, \Delta_i\Big]. \tag{207}$$

Because (206) is imposed at every candidate, $S(r^a)$ equals the gap between the two margins of the combined condition (61), and its zero is the interior stationary optimum. S is increasing in r^a in the neighborhood of the root; a bracket walk plus bisection delivers r^{a*}, and Figure 10 plots the two margins whose crossing it computes.

E.3 Transition dynamics: the time-varying discretionary system

Consider an unanticipated date-0 innovation to an exogenous aggregate path $\mathbf{Z} = \{Z_t\}$. The economy starts at the pre-shock stationary optimum and the discretionary planner re-optimizes from the inherited state. The transition system is the optimality system of Proposition 5 with $\psi \equiv 0$, evaluated date by date on a uniform grid $t \in \{0, \Delta t, \ldots, T_n \Delta t\}$ and closed by the requirement that the economy is at the (post-shock) stationary optimum from the truncation horizon onward. The baseline uses a monthly step. The solve horizon requires care: the planner's co-states spread the inflation levy over the economy's slow fiscal and distributional adjustments, so finite-horizon paths carry truncation error long before the terminal date; reported paths are computed at several horizons and extrapolated, as described in Section E.4. Transitions are computed at the baseline remittance limit $\bar{R} = 0$, so $R_t = 0$ along the path. The system stacks the unknown sequences $\mathbf{U}$ against an equal number of residual conditions, written $\mathbf{R}(\mathbf{U}; \mathbf{Z}) = \mathbf{0}$; the blocks below compute one evaluation of the map, and Figure 14 displays its structure. Section E.4 forms the sequence-space Jacobian of this stacked system.

Unknowns. The aggregate unknown sequences are

$$\mathbf{U} = \{r^a_t,\; i^b_t,\; \vartheta_t,\; Q_t\}_{t=0}^{T_n - 1}, \tag{208}$$

$4T_n$ in total. The discretionary revaluation is the implied jump $\nu \equiv Q_0/Q_{ss}$: in detrended units a jump in the initial price level leaves the inherited distribution unchanged and rescales real balances and the initial real stocks, $(A^{b,T}_0, A^{b,CB}_0) = \nu\,(A^{b,T}_{ss}, A^{b,CB}_{ss})$. Everything else is computed from $\mathbf{U}$ by the blocks below.

Fiscal block. Treasury debt runs forward from $A^{b,T}_0 = \nu A^{b,T}_{ss}$,

$$A^{b,T}_{t+\Delta t} = A^{b,T}_t + \Delta t \Big[\bar{T} + \kappa \bar{b} - \bar{R} - (\kappa - r^b_t)\, A^{b,T}_t\Big], \qquad T_t = \bar{T} + \kappa(\bar{b} - A^{b,T}_t), \tag{209}$$

with $r_t^b = r_t^a + i_t^b$. The block's output is the Treasury debt path $\{A_t^{b,T}\}$ and the transfer path $\{T_t\}$.

Household block. Given $\{r_t^a, i_t^b, T_t\}$, the candidate real-balance path $\{Q_t\}$, and the shock path: solve the HJB (200) backward from the terminal condition $\mathbf{W}_{T_n} = \mathbf{W}_{ss}$, extracting policies, thresholds, and generators $\mathbf{A}_t$ at each date; run the KFE forward from g_0; integrate for money demand $Q_t^{dem} = \sum_i y_{i,t}\, g_{i,t}\, \Delta_i$, total wealth $A_t = \sum_i \omega_i\, g_{i,t}\, \Delta_i$, and household bond demand $A_t^{b,hh} = A_t - Q_t^{dem}$. The block's output is the policies and generators $\{\mathbf{A}_t\}$, the distribution path $\{g_t\}$, and the aggregates $\{Q_t^{dem}, A_t, A_t^{b,hh}\}$.

Central-Bank block. Real balances are the unknown path itself, $A_t^a = Q_t$, with clearing imposed in the residual stack; the implied money growth $g_t^M = (Q_{t+\Delta t} - Q_t)/(Q_t \Delta t) - r_t^a$ is post-processing and enters no residual. The central bank's holdings run forward from $A_0^{b,CB} = \nu A_{ss}^{b,CB}$ via its flow budget in real terms,

$$A_{t+\Delta t}^{b,CB} = A_t^{b,CB} + \Delta t \Big[\frac{Q_{t+\Delta t} - Q_t}{\Delta t} + r_t^b A_t^{b,CB} - r_t^a Q_t - \bar{R} \Big]. \tag{210}$$

The block's output is the central bank's position path $\{A_t^{b,CB}\}$; it enters bond-market clearing through the float $A_t^{b,T} - A_t^{b,CB}$ and weights the shadow price ϑ_t in the two optimality conditions.

Planner co-state block. Given $\{\vartheta_t\}$ and the household block outputs:

(i) The market-clearing shadow values are algebraic in the ϑ-path, from the discrete co-state equations (55) and (56):

$$\mu_t = (\rho - r_t^a)\, \vartheta_t - \frac{\vartheta_{t+\Delta t} - \vartheta_t}{\Delta t}, \qquad \varphi_t = \mu_t - i_t^b\, \vartheta_t, \tag{211}$$

with $\vartheta_{T_n} = \vartheta_{ss}$.

(ii) The social value runs backward with the same factorized operator as the HJB,

$$\Big[(\rho + \frac{1}{\Delta t})\mathbf{I} - \mathbf{A}_t \Big] \gamma_t = \boldsymbol{\mathcal{U}}_t + \mu_t\, \mathbf{y}_t + \varphi_t(\boldsymbol{\omega} - \mathbf{y}_t) + \frac{1}{\Delta t}\, \gamma_{t+\Delta t}, \tag{212}$$

from $\gamma_{T_n} = \gamma_{ss}$. The uniform source $-\zeta_t T_t$ is omitted: it shifts γ_t by a time function that drops from every gradient moment below.

(iii) The average social value is $\theta_t = \sum_i \partial_\omega \gamma_{i,t}\, g_{i,t}\, \Delta_i$, and the Treasury's shadow price runs backward from $\zeta_{T_n} = \zeta_{ss}$ via the discrete form of (57),

$$\zeta_t = \frac{\zeta_{t+\Delta t} + \Delta t\, (\varphi_t + \kappa\, \theta_t)}{1 + \Delta t\, (\rho - r_t^b + \kappa)}. \tag{213}$$

(iv) The wedges and multipliers follow from (201) with ϑ_t in the portfolio term.

The block's output is the gradient moments of γ_t, the Treasury shadow price path $\{\zeta_t\}$, and the wedges $\{\eta_{i,t}\}$; together they assemble the two optimality-condition residuals below.

Residuals. The stacked system $\mathbf{R}(\mathbf{U};\mathbf{Z}) = \mathbf{0}$ imposes, at each date $t = 0, \ldots, T_n - 1$:

1. **Bond-market clearing:**

$$R_t^{bmc} \equiv A_t^{b,hh} - (A_t^{b,T} - A_t^{b,CB}) = 0. \tag{214}$$

2. **Optimal inflation):** the discrete version of (59),

$$R_t^{r^a} \equiv \sum_i \omega_i \, \partial_\omega \gamma_{i,t} \, g_{i,t} \, \Delta_i - \left[\vartheta_t A_t^a + \zeta_t A_t^{b,T} - \vartheta_t A_t^{b,CB}\right] = 0. \tag{215}$$

3. **Optimal bond rate:** the discrete version of (58),

$$R_t^{i^b} \equiv \sum_i (\omega_i - y_{i,t}) \, \partial_\omega \gamma_{i,t} \, g_{i,t} \, \Delta_i - \left[\zeta_t A_t^{b,T} - \vartheta_t A_t^{b,CB} + \sum_{i:\, \omega_i > \underline{\omega}_t(z)} \eta_{i,t} \, u'(c_{i,t}) \, \Delta_i\right] = 0. \tag{216}$$

4. **Money-market clearing**:

$$R_t^{mc} \equiv Q_t^{dem} - Q_t = 0. \tag{217}$$

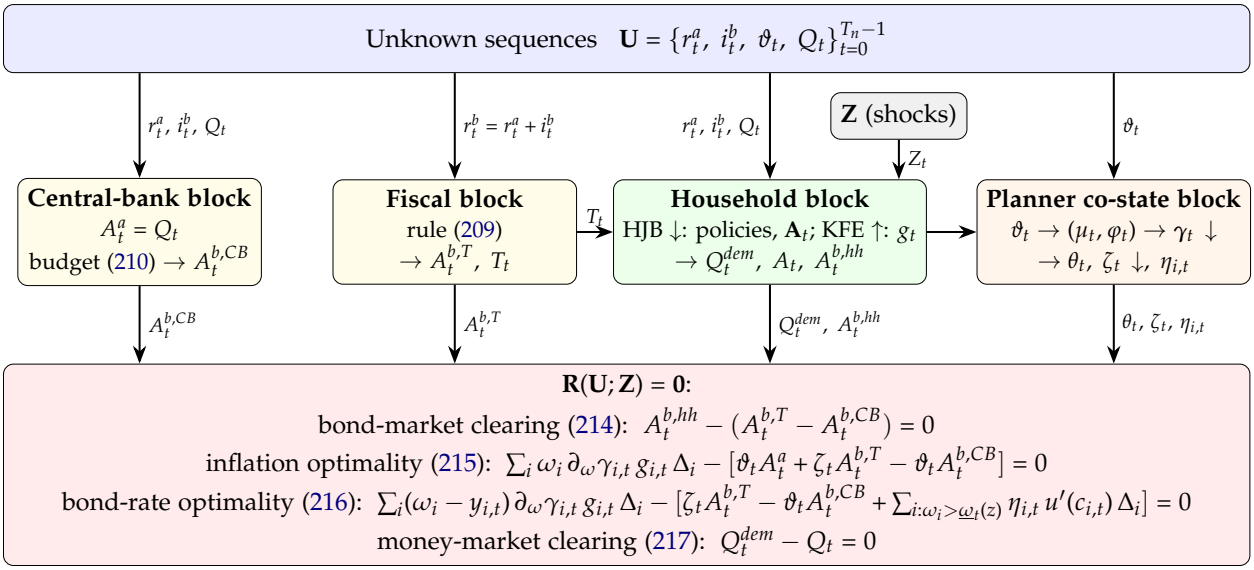


Figure 14: The directed acyclic graph of detrended, separated-discretion optimality system.

E.4 Sequence-space Jacobian and solution algorithm

The residuals depend on $\mathbf{U}$ directly through the budget and identity terms and indirectly through the block outputs. By the chain rule,

$$\mathbf{J} = \left.\frac{\partial \mathbf{R}}{\partial \mathbf{U}}\right|_{\text{blocks fixed}} + \frac{\partial \mathbf{R}}{\partial\, \text{blocks}} \cdot \frac{\partial\, \text{blocks}}{\partial \mathbf{U}}. \tag{218}$$

The direct terms and the fiscal- and central-bank-block paths are analytic (triangular recursions in $\mathbf{U}$). The household and planner co-state blocks are differentiated by linearized sweeps. For the household block, the policies are differentiated in closed form through their first-order conditions, which yields the perturbed drifts, expenditure policies, and generators. The linearized KFE then runs forward, giving the response of the distribution path, and the perturbed aggregates $\{Q_t^{dem}, A_t, A_t^{b,hh}\}$ follow by integration. This is the sequence-space Jacobian method of Auclert et al. (2021), computed as one linearized backward–forward pass per unit direction. For the planner's co-states, the same pass in the household block differentiates the γ recursion explicitly. The KFE (212) is solved backward with the same operator, with three sources: the perturbed loadings $(d\mu_t, d\varphi_t)$, analytic in the ϑ direction by (211); the perturbed flow utility and expenditure policies from the household pass; and the stationary γ acted on by the perturbed generator. The perturbed gradient moments then combine three channels, the perturbation of γ itself, the perturbation of the policies entering the moment weights, and the response of the distribution from the forward pass. The wedge integral in (216) is differentiated the same way, with its ϑ_t loading analytic from (201) and the movement of the full-depletion threshold included. All Jacobians are evaluated once at the stationary optimum and assembled into $\mathbf{J}_{ss}$.

The reported impulse responses to transitory shocks, whose paths return to the initial stationary optimum, are linear. One linearized pass along the shock path with the unknowns held at their stationary values gives the direct effect of the shock on the residual stack, $(\partial \mathbf{R}/\partial \mathbf{Z})\, d\mathbf{Z}$, and the response of the unknowns solves $\mathbf{J}_{ss}\, d\mathbf{U} = -(\partial \mathbf{R}/\partial \mathbf{Z})\, d\mathbf{Z}$.[44] Reported paths are the stationary values plus $d\mathbf{U}$. Larger shocks are solved nonlinearly: initialize $\mathbf{U}^0$ at the stationary values, evaluate the blocks and residuals, and take damped Newton steps through the same filtered inverse with the Jacobian refreshed at the shocked path; reported transitions there are differences between the shocked and zero-shock solves at the same settings, and departures from proportional scaling in the shock size measure genuine nonlinearity rather than numerical error.

[44] For some impulse responses, we use the singular value decomposition of $\mathbf{J}_{ss}$ with each direction damped by the Tikhonov filter $s/(s^2 + \lambda^2)$ — a smooth filter. The filter matters only along a cluster of near-degenerate directions that reshuffle the timing of inflation at nearly fixed real paths; it selects the minimal-norm member of that family, and the shocks studied here load on it negligibly.

F Model with Private Capital

G Model with Portfolio Adjustment Costs

This appendix collects the alternative formulation of the model in which the household's portfolio rebalancing between liquid (money) and illiquid (bond) accounts is modeled as a continuous flow d_t subject to a strictly convex adjustment cost $\Phi(d_t)$, following Kaplan and Violante (2014) and Kaplan et al. (2018). This formulation differs from the baseline environment in its treatment of the household's portfolio margin but coincides on most of the economic intuition. Each subsection mirrors a section of the main body. Propositions are stated with their proofs placed inline immediately after each statement.

G.1 Environment with adjustment costs

This subsection restates the household and government environment with continuous-flow portfolio adjustment costs $\Phi(d_t)$. The aggregate preferences and the government budget structure are unchanged from the main body and are not repeated; we re-state them only where the adjustment-cost formulation alters the equations. The state space has three dimensions $x = (a, b, z)$: liquid wealth a, illiquid wealth b, and idiosyncratic productivity z.

Assets and household budget constraints. From the standpoint of households, the assets differ in two ways: liquidity for lumpy expenditures and in the cost of portfolio adjustment. In the event of a lumpy expenditure shock, households can only use money. Additionally, transferring wealth between liquid and illiquid accounts incurs an adjustment cost $\Phi(d_t)$, where d_t is the flow deposit rate. The function Φ satisfies $\Phi(0) = 0$, $\Phi'(0) = 0$, and $\Phi''(d) > 0$ for all d: zero cost and zero marginal cost at $d = 0$, with strictly increasing marginal cost for deposits or withdrawals.

The household's law of motion for liquid wealth (outside of lumpy expenditure shocks) is

$$\dot{a}_t = r_t^a a_t + z_t - c_t + \mathcal{T}_t(a, b, z) - d_t - \Phi(d_t), \tag{219}$$

where $r_t^a = \dot{\phi}_t/\phi_t$ is the real return on liquid wealth. The adjustment cost $\Phi(d_t)$ is paid out of liquid wealth. They face a no-borrowing constraint $a_t \geq 0$. The law of motion for illiquid wealth is

$$\dot{b}_t = r_t^b b_t + d_t, \tag{220}$$

with a borrowing constraint $b_t \geq \underline{b}$, where $\underline{b} \leq 0$, and the real return is $r_t^b = i_t^b + \dot{\phi}_t/\phi_t$. The household would like to, in the event of a lumpy-expenditure shock, discontinuously withdraw from

their illiquid account, but this is not allowed—deposits are continuous flows and the adjustment cost prevents instantaneous rebalancing.

Decentralized equilibrium with adjustment costs. Households take policy and prices (ϕ_t, i_t^b) as given. Let $W_t(a, b, z)$ represent a household's value function at time t. It satisfies the Hamilton-Jacobi-Bellman equation

$$\rho W_t(a,b,z) = \max_{c_t, y_t, d_t} u(c_t) + \alpha U(y_t) + \tilde{\mathcal{A}}_t(a,b,z,c,y,d)\left[W_t\right] + \partial_t W_t, \tag{221}$$

where the constraints on choices are $c_t \geq 0$, $0 \leq y_t \leq a_t$, and d_t subject to the state constraints $a_t \geq 0$ and $b_t \geq \underline{b}$. The operator $\tilde{\mathcal{A}}_t$ represents the infinitesimal generator of the idiosyncratic state (a, b, z) and is given by

$$\begin{aligned}\tilde{\mathcal{A}}_t[W] = &\left[r_t^a a + z_t - c + \mathcal{T}_t - d_t - \Phi(d_t)\right]\partial_a W + \left[r_t^b b + d_t\right]\partial_b W \\ &+ \alpha\left[W(a-y,b,z) - W(a,b,z)\right] + \mathcal{A}_z W(a,b,z).\end{aligned} \tag{222}$$

The generator has four components: continuous changes in liquid wealth (first term), continuous changes in illiquid wealth (second term), discrete changes in liquid wealth from lumpy consumption (third term), and idiosyncratic productivity transitions (fourth term).

The optimization problem leads to three optimality conditions:

$$u'(c_t) = \partial_a W_t, \tag{223}$$

$$U'(y_t) \geq \partial_a W_t(a - y_t, b, z), \tag{224}$$

$$(1 + \Phi'(d_t))\,\partial_a W_t = \partial_b W_t. \tag{225}$$

Equation (223) equates the marginal utility of flow consumption with the marginal value of liquid wealth. Equation (224) holds with equality if $y_t < a$; it governs lumpy consumption, which can only be financed with liquid wealth. Equation (225) equates the marginal cost of depositing one dollar (the dollar plus its adjustment cost, valued at $\partial_a W_t$) with the marginal value of illiquid wealth $\partial_b W_t$. The strict convexity of Φ ensures a unique interior solution for d_t at every state.

We can re-write the HJB plugging in the household policies as

$$\rho W_t = u(c_t) + \alpha U(y_t) + \mathcal{A}_t\left[W_t\right] + \partial_t W_t, \tag{226}$$

where $\mathcal{A}_t = \tilde{\mathcal{A}}_t(a, b, z, c_t, y_t, d_t)$. The equilibrium measure G_t over household states (a, b, z) satisfies

the Kolmogorov forward equation:

$$\partial_t G_t = \mathcal{A}_t^*[G_t], \tag{227}$$

where $\mathcal{A}_t^*$ is the adjoint of the generator $\mathcal{A}_t$. The stationary distribution may contain mass points at the borrowing constraints—at $a = 0$ (for each (b, z)), at $b = \underline{b}$ (for each (a, z)), or at the double corner $(0, \underline{b})$ (for each z). We denote the mass at $a = 0$ by $m_t(b, z)$, the mass at $b = \underline{b}$ by $\tilde{m}_t(a, z)$, and the absolutely continuous part of G_t by its density $g_t(a, b, z)$. The adjoint $\mathcal{A}_t^*$ acts as a differential operator on the density g_t; separate boundary ODEs govern the evolution of the mass points. Market clearing requires

$$\int a \, dG_t(a, b, z) = \phi_t M_t, \tag{228}$$

$$\int b \, dG_t(a, b, z) = \phi_t B_t. \tag{229}$$

The decentralized competitive equilibrium under adjustment costs takes the same form as Definition 1 in the main body, with the household HJB (226), KFE (227), and market-clearing conditions (228)–(229) replacing their no-adjustment-cost counterparts. The initial price level ϕ_0 is uniquely pinned down by the same fiscal-anchor argument, restricted to non-explosive equilibria.

Characterization of the decentralized equilibrium with adjustment costs. We re-establish some properties of the decentralized equilibrium, now with adjustment costs. We impose Assumption 3 on the transfer rule that guarantees all households can finance a strictly positive consumption flow.

Assumption 3 (Feasibility with adjustment costs)**.** *For all t, transfers and the borrowing limit satisfy $\underline{z}_t + \mathcal{T}_t(\underline{b}, \underline{z}) + r_t^b \, \underline{b} - \Phi(-r_t^b \, \underline{b}) > 0$, for any $0 \leq r_t^b \leq \rho$.*

Assumption 3 rules out confiscatory transfers that would leave a household at $\dot{a} \leq 0$ with non-positive resources.

Proposition 10 (Value function properties and boundary conditions with adjustment costs)**.** *Under Assumption 3:*

i) Viscosity solution: $W_t(a, b, z)$ is a constrained viscosity solution of the HJB equation (226) on $[0, \infty) \times [0, \infty) \times \mathcal{Z}$.

ii) Boundary conditions. For each t, b, and $z \in \mathcal{Z}$, $\partial_a W_t(0, b, z) < \infty$. At $a = 0$, the consumption Euler $u'(c_t) = \partial_a W_t$ holds and the state-constraint boundary condition satisfies

$$\partial_a W_t(0, b, z) \geq u'\left(z_t + \mathcal{T}_t(0, b, z) - d_t(0, b, z) - \Phi(d_t(0, b, z))\right), \tag{230}$$

with equality iff $\dot{a}_t(0, b, z) = 0$. *At* $b = \underline{b}$, *the state constraint* $\dot{b}_t \geq 0$ *requires* $d_t \geq -r_t^b \underline{b}$. *At the double corner* ($a = 0$, $b = \underline{b}$), *both state constraints are active and the consumption bound satisfies*

$$\partial_a W_t(0, \underline{b}, z) \geq u'\left(z_t + \mathcal{T}_t + r_t^b \underline{b} - \Phi(-r_t^b \underline{b})\right), \tag{231}$$

with equality iff $\dot{a}_t = \dot{b}_t = 0$.

iii) Monotonicity and concavity: W_t *is strictly increasing and strictly concave in* a *for each* (b, z), *and strictly increasing and strictly concave in* b *for each* (a, z).

iv) Zero savings at the double corner: At the double corner $(a, b) = (0, \underline{b})$ *with the lowest productivity type* $\underline{z}$, $\dot{a}_t(0, \underline{b}, \underline{z}) = \dot{b}_t(0, \underline{b}, \underline{z}) = 0$.

Proof. We establish the viscosity solution property, boundary conditions, concavity, and smooth pasting on the state space $[0, \infty) \times [0, \infty) \times \mathcal{Z}$. Define the HJB operator $\mathcal{H}_t : [0, \infty) \times [0, \infty) \times \mathcal{Z} \times \mathbb{R}^2 \to \mathbb{R}$ by

$$\begin{aligned} \mathcal{H}_t(a, b, z, p_a, p_b) \equiv\ & \rho W_t - u(c_t) - \alpha U(y_t) \\ & - \left[r_t^a a + z_t - c_t + \mathcal{T}_t - d_t - \Phi(d_t)\right] p_a - \left[r_t^b b + d_t\right] p_b \\ & - \alpha\left[W_t(a - y_t, b, z) - W_t(a, b, z)\right] - \mathcal{A}_z W_t - \partial_t W_t, \end{aligned} \tag{232}$$

so that the HJB equation (226) reads $\mathcal{H}_t(a, b, z, \partial_a W_t, \partial_b W_t) = 0$.

A continuous function v is a *constrained viscosity subsolution* of (226) if, for each $(a_0, b_0, z_0) \in (0, \infty) \times (0, \infty) \times \mathcal{Z}$ and each smooth ϕ such that $v - \phi$ attains a local maximum at (a_0, b_0, z_0) with $v(a_0, b_0, z_0) = \phi(a_0, b_0, z_0)$:

$$\mathcal{H}_t(a_0, b_0, z_0, \partial_a \phi(a_0, b_0), \partial_b \phi(a_0, b_0)) \leq 0. \tag{233}$$

A continuous function v is a *constrained viscosity supersolution* of (226) if, for each $(a_0, b_0, z_0) \in [0, \infty) \times [0, \infty) \times \mathcal{Z}$ (including the boundaries $a_0 = 0$ and $b_0 = \underline{b}$) and each smooth ϕ such that $v - \phi$ attains a local minimum at (a_0, b_0, z_0) with $v(a_0, b_0, z_0) = \phi(a_0, b_0, z_0)$:

$$\mathcal{H}_t(a_0, b_0, z_0, \partial_a \phi(a_0, b_0), \partial_b \phi(a_0, b_0)) \geq 0. \tag{234}$$

Part (i): W_t **is a constrained viscosity solution.** Fix $(a_0, b_0, z_0) \in (0, \infty) \times (0, \infty) \times \mathcal{Z}$ and let ϕ be a smooth test function such that $W_t - \phi$ attains a local maximum at (a_0, b_0, z_0) with $W_t(a_0, b_0, z_0) =$

$\phi(a_0, b_0, z_0)$. By the Dynamic Programming Principle applied to W_t,[45] for any $h > 0$,

$$W_t(a_0, b_0, z_0) = \mathbb{E}\left[\int_t^{t+h} e^{-\rho(s-t)}\mathcal{U}_s^*\, ds + e^{-\rho h} W_{t+h}(a_{t+h}^*, b_{t+h}^*, z_{t+h})\right].$$

Applying Dynkin's formula to $e^{-\rho s}\phi(a_s, b_s, z_s)$ over $[t, t+h]$, using the local maximum property of ϕ on continuous and Poisson-jump segments, dividing by h, and sending $h \to 0$ yields the subsolution condition. The supersolution argument is symmetric, using the DPP inequality $W_t \geq \mathbb{E}[\cdots]$ for any admissible control and applying Dynkin's formula to ϕ under $W_t \geq \phi$ throughout; this extends to $a_0 = 0$ and $b_0 = \underline{b}$ by restricting the admissible control set to policies that maintain the state constraints.

Part (ii): Boundary conditions. At $(0, b, z)$, Assumption 3 guarantees positive resources and the Inada condition $u'(0) = \infty$ ensures $c_t(0, b, z) > 0$. The Euler condition $\partial_a W_t(0, b, z) = u'(c_t(0, b, z)) < u'(0) = \infty$ gives finiteness. At $a = 0$, the constraint $y_t \leq a = 0$ forces $y_t(0, b, z) = 0$ and the no-borrowing constraint $\dot{a} \geq 0$ requires $c_t \leq z_t + \mathcal{T}_t - d_t - \Phi(d_t)$. Since $u'' < 0$, $\partial_a W_t(0, b, z) = u'(c_t) \geq u'(z_t + \mathcal{T}_t - d_t - \Phi(d_t))$, establishing (230). At $b = \underline{b}$, the state constraint $\dot{b}_t = r_t^b \underline{b} + d_t \geq 0$ directly requires $d_t \geq -r_t^b \underline{b}$. At the double corner, $d_t = -r_t^b \underline{b}$ and $c_t = z_t + \mathcal{T}_t + r_t^b \underline{b} - \Phi(-r_t^b \underline{b})$ when both constraints bind; Assumption 3 guarantees this is positive, and the Euler equation gives (231).

Part (iii): Monotonicity and concavity. Fix t, (b, z), $\lambda \in (0, 1)$, and $a^0 \neq a^1$ with $\hat{a} = \lambda a^0 + (1-\lambda)a^1$. Let $\pi^i = (c_s^i, y_s^i, d_s^i)_{s \geq t}$ denote the optimal plan from (a^i, b, z) and define the mixture plan $\hat{\pi} \equiv \lambda\pi^0 + (1-\lambda)\pi^1$. Since the liquid budget (219) is affine in (a, c, d) for given (b, z), the mixture liquid wealth $\hat{a}_s \equiv \lambda a_s^0 + (1-\lambda)a_s^1$ satisfies the budget constraint under $\hat{\pi}$ (the convexity of Φ implies the mixture wastes weakly more in adjustment costs, preserving feasibility). The illiquid budget is affine in (b, d), so $\hat{b}_s$ is also feasible. By strict concavity of u and U, the mixture plan yields strictly higher utility than the convex combination of π^0 and π^1 on a set of positive measure. The analogous argument exchanging a and b gives strict concavity in b. Strict monotonicity in each argument follows from the strict relaxation of the budget constraint by additional wealth.

Part (iv): Zero savings at the double corner for $\underline{z}$. At the double corner with the lowest productivity, $y_t = 0$, $d_t = -r_t^b \underline{b}$, and $\dot{a}_t = z_t - c_t + \mathcal{T}_t + r_t^b \underline{b} - \Phi(-r_t^b \underline{b})$. The same phase-diagram argument as in Achdou et al. (2022, Section 4.1) applies: if $\dot{a}_t > 0$ at $(0, \underline{b}, \underline{z})$, the coupled ODE system for $(\partial_a W_t, \dot{a}_t)$ near $a = 0$ lies on an unstable manifold that diverges, inconsistent with a globally well-defined value function. The unique consistent solution requires $\dot{a}_t(0, \underline{b}, \underline{z}) = 0$, which combined with $d_t = -r_t^b \underline{b}$ implies $\dot{b}_t = 0$. □ □

Proposition 10 extends the results in Rocheteau et al. (2018a) and Achdou et al. (2022) to non-

[45] Invoking the DPP and Dynkin's formula requires that optimal drifts $|\dot{a}_s^*|$ and $|\dot{b}_s^*|$ are bounded along optimal paths. The liquid drift is bounded because c^* and d^* are finite (the last ensured by the strictly convex adjustment cost Φ). The illiquid drift $r^b b + d^*$ is bounded because d^* is finite. These extend the regularity results of Rocheteau et al. (2018a, Lemmas 1–2) to the two-asset setting.

stationary environments with idiosyncratic productivity risk and lumpy consumption, which allows the use of (221) to characterize W_t uniquely.

Proposition 11 (Lumpy consumption with adjustment costs)**.** *For each $t \geq 0$, $b \geq \underline{b}$, and $z \in \mathcal{Z}$: i) there exists a threshold $\bar{a}_t(b,z) > 0$ such that $y_t(a,b,z) = a$ for all $a \in (0, \bar{a}_t(b,z)]$, where the threshold satisfies $U'(\bar{a}_t(b,z)) = \partial_a W_t(0,b,z)$, and ii) for each $a > 0$, $0 \leq \partial_a y_t(a,b,z) \leq 1$.*

Proof. At state (a,b,z), the optimal lumpy consumption solves $\max_{0 \leq y \leq a}\{U(y) + W_t(a-y,b,z)\}$. Since U is strictly concave and $W_t(\cdot,b,z)$ is strictly concave in a (Proposition 10(iii)), the objective is strictly concave in y and the maximizer is unique.

Part (i): Full depletion. The corner $y = a$ is optimal iff $U'(a) \geq \partial_a W_t(0,b,z)$. By the Inada condition $U'(0) = \infty$ and the finiteness $\partial_a W_t(0,b,z) < \infty$ (Proposition 10(ii)), the intermediate value theorem guarantees a unique $\bar{a}_t(b,z) > 0$ solving $U'(\bar{a}_t(b,z)) = \partial_a W_t(0,b,z)$. Uniqueness follows from $U'' < 0$. Fix $a \in (0, \bar{a}_t(b,z)]$ and suppose for contradiction that $y_t(a,b,z) < a$. The Euler equation (224) would then hold with equality: $U'(y_t) = \partial_a W_t(a - y_t, b, z) < \partial_a W_t(0,b,z) = U'(\bar{a}_t(b,z))$ since $\partial_a W_t(\cdot,b,z)$ is strictly decreasing. By $U'' < 0$, this gives $y_t > \bar{a}_t(b,z) \geq a$, contradicting $y_t < a$. Hence $y_t(a,b,z) = a$ for all $a \in (0, \bar{a}_t(b,z)]$.

Part (ii): Monotonicity. For $a \in (0, \bar{a}_t(b,z)]$, $y_t(a,b,z) = a$, so $\partial_a y_t = 1$. For $a > \bar{a}_t(b,z)$, the interior FOC $U'(y_t) = \partial_a W_t(a - y_t, b, z)$ holds. Differentiating with respect to a:

$$\partial_a y_t = \frac{\partial_{aa} W_t(a - y_t, b, z)}{U''(y_t) + \partial_{aa} W_t(a - y_t, b, z)} \in (0,1),$$

where both numerator and denominator are negative. □ □

Proposition 11 establishes that households near the liquid borrowing constraint fully deplete their wealth in the event of a lumpy consumption shock.

G.2 Optimal consolidated policy with adjustment costs

The consolidated planner chooses $\{\dot{M}_s, \dot{B}_s, \mathcal{T}_s\}$ to maximize utilitarian welfare subject to the adjustment-cost competitive equilibrium (226)–(229). The optimality system using the same approach as in Appendix C. The full derivation, together with the four steady-state characterizations below, is collected in a companion note available upon request. It introduces the social value $\gamma_s(x)$ (the co-state on the distribution), the promise multiplier $\psi_s(x)$ (the co-state on the household value, governing commitment), the shadow price of government revenue θ_s, and three household-implementability wedges on the consumption, lumpy, and the (new) deposit FOCs (223)–(225),

$$\tilde{\chi}_s \equiv \partial_a \gamma_s - u'(c_s), \qquad \tilde{\eta}_s \equiv \partial_a \gamma_s(a - y_s, b, z) - U'(y_s), \qquad \tilde{\delta}_s \equiv (1 + \Phi'(d_s))\,\partial_a \gamma_s - \partial_b \gamma_s. \tag{235}$$

The deposit wedge $\tilde{\delta}_s$ is new relative to the baseline: it is the planner's counterpart of the household's deposit FOC residual, generically nonzero, and vanishes when the planner equates the social marginal values of liquid and illiquid wealth. The lumpy wedge $\tilde{\eta}_s$ carries no liquidity-premium term: unlike the baseline, where money is held only to finance lumpy purchases so the foregone interest enters the lumpy margin, here liquid wealth is a smooth portfolio choice and the liquidity premium is priced on the deposit margin. Under lump-sum transfers the shadow price satisfies

$$\theta_s = \int \partial_a \gamma_s \, dG_s + \int \psi_s \, \partial_a W_s \, dx, \tag{236}$$

and the inflation and bond-rate margins are governed by the wealth-weighted social value

$$R_s \equiv \int [a \, \partial_a \gamma_s + b \, \partial_b \gamma_s] \, dG_s + \int [a \, \partial_a W_s + b \, \partial_b W_s] \, \psi_s \, dx, \tag{237}$$

through the static revaluation $R_s = \theta_s A_s$ (with $A_s = \phi_s(M_s + B_s) = A_s^a + A_s^b$ total real liabilities) and the bond-rate condition $\int b \, \partial_b \gamma_s \, dG_s + \int b \, \psi_s \partial_b W_s \, dx = \theta_s A_s^b$. The forward equation for ψ_s carries an additional source term from the deposit constraint, $-\partial_a[\delta_s(1 + \Phi'(d_s))] + \partial_b \delta_s$ with $\delta_s = \tilde{\delta}_s \, g_s / (\Phi''(d_s) \partial_a W_s)$, alongside the consumption and lumpy sources.

Unrestricted transfers. When transfers can condition on (a, b, z), the planner equalizes the adjusted social marginal value of wealth across agents, the promise multiplier vanishes for both $\lambda \to 0$ and $\lambda \to \infty$, and the deposit wedge $\tilde{\delta}_s$ vanishes as well.

Proposition 12 (Friedman rule under unrestricted transfers, adjustment costs)**.** *Under unrestricted transfers and either full discretion or full commitment, the optimal policy in the adjustment-cost model sets $B_t = 0$ and $M_t > 0$ with $r_t^a = \rho$ (Friedman rule), $i_t^b = 0$, and $d_t = 0$ for all households. The allocation is time and state invariant and coincides with the first-best in (2)–(3).*

At the Friedman rule bonds are redundant: with $r_t^a = r_t^b = \rho$ they offer no return advantage over money but lack its liquidity, so any $B_t > 0$ would induce withdrawals at a deadweight adjustment cost $\Phi(d) > 0$. Setting $B_t = 0$ eliminates the illiquid account ($b_t = 0$, $d_t = 0$, $\Phi(0) = 0$), and the model collapses to the pure-money first best of the main body. The transactions motive keeps the wealth distribution non-degenerate, since the planner cannot replicate the discrete lumpy purchase with transfers.

Assumption 4 (Bounded relative risk aversion)**.** *There exists $\bar{\gamma} < \infty$ such that $-cu''(c)/u'(c) \leq \bar{\gamma}$ for all $c > 0$, and absolute risk aversion $-u''(c)/u'(c)$ remains finite as $c \downarrow 0$.*

Discretion with lump-sum transfers. Under full discretion ($\psi_s \equiv 0$), $\theta_{ss} = \int \partial_a \gamma_{ss} \, dG_{ss}$ and the two margins separate cleanly across the assets: subtracting the bond-rate condition from

$R_{ss} = \theta_{ss} A_{ss}$ isolates the liquid margin.

Proposition 13 (Optimal consolidated policy under full discretion, adjustment costs)**.** *Under Assumptions 3 and 4 and full discretion with lump-sum transfers, the optimal steady state satisfies $r_{ss}^{a*} < \rho$ and is characterized by the liquid-wealth and illiquid-wealth margins*

$$\int a\, \partial_a \gamma_{ss}\, dG_{ss} = \theta_{ss} A_{ss}^a, \qquad \int b\, \partial_b \gamma_{ss}\, dG_{ss} = \theta_{ss} A_{ss}^b, \tag{238}$$

with $\theta_{ss} = \int \partial_a \gamma_{ss}\, dG_{ss}$. The illiquid margin is the bond-rate (i^b) first-order condition; the inflation (δg^M) condition is their sum, $R_{ss} = \theta_{ss} A_{ss}$. Equivalently, in covariance form,

$$\mathrm{Cov}_{G_{ss}}(a, \partial_a \gamma_{ss}) = 0, \qquad \mathrm{Cov}_{G_{ss}}(b, \partial_b \gamma_{ss}) = -A_{ss}^b \int (\partial_b \gamma_{ss} - \partial_a \gamma_{ss}) dG_{ss}. \tag{239}$$

The liquid margin sets the inflation tax so that liquid wealth and its social marginal value are uncorrelated, exactly as in the baseline. The illiquid margin sets the liquidity premium: its right-hand side equals $-A_{ss}^b \int [\Phi'(d_{ss}) \partial_a \gamma_{ss} - \tilde{\delta}_{ss}]\, dG_{ss}$, the average marginal adjustment-cost dissipation net of the average deposit wedge. When $\tilde{\delta}_{ss} > 0$ on average, the planner tilts the balance sheet toward money to correct the household's undersupply of liquidity. The liquidity premium does not vanish as $\Phi \to 0$: money remains the only lumpy-eligible asset and earns the lower return r_{ss}^a, so households still hold a costly liquid buffer and the friction does not disappear. The adjustment cost is an alternative pricing of this money-demand friction—through the cost of rebalancing rather than the baseline's requirement that money be pre-held—and the $\Phi \to 0$ limit is singular, approaching the baseline lumpy economy, where the same premium is carried by the bond-rate condition's implementability term $\int \eta_{ss} u'(c_{ss})\, dx$ instead of the deposit margin.

Commitment with lump-sum transfers. Under full commitment ($\lambda \to 0$), $\psi_{ss} \neq 0$ and the planner internalizes how policy shapes the distribution future planners inherit, so $\theta_{ss} = \int \partial_a \gamma_{ss}\, dG_{ss} + \Psi_{ss}$ with commitment premium $\Psi_{ss} \equiv \int \psi_{ss}\, u'(c_{ss})\, dx$.

Proposition 14 (Optimal consolidated policy under full commitment, adjustment costs)**.** *Under Assumptions 3 and 4 and full commitment with lump-sum transfers, the stationary equilibrium satisfies $r_{ss}^{a,\mathrm{comm}} < \rho$ and the liquid-wealth and illiquid-wealth margins*

$$\int [a\, \partial_a \gamma_{ss}\, g_{ss} + a\, \psi_{ss} \partial_a W_{ss}]\, dx = \theta_{ss} A_{ss}^a, \qquad \int [b\, \partial_b \gamma_{ss}\, g_{ss} + b\, \psi_{ss} \partial_b W_{ss}]\, dx = \theta_{ss} A_{ss}^b. \tag{240}$$

The illiquid margin is again the bond-rate condition, the two summing to the inflation condition $R_{ss} =$

$\theta_{ss} A_{ss}$. *Equivalently, in covariance form,*

$$\mathrm{Cov}_{G_{ss}}(a, \partial_a \gamma_{ss}) + \Psi^a_{ss} = 0, \tag{241}$$

$$\mathrm{Cov}_{G_{ss}}(b, \partial_b \gamma_{ss}) + \Psi^b_{ss} = -A^b_{ss} \int (\partial_b \gamma_{ss} - \partial_a \gamma_{ss}) dG_{ss} - A^b_{ss} \int \psi_{ss}\, \Phi'(d_{ss})\, u'(c_{ss})\, dx, \tag{242}$$

where $\Psi^a_{ss} \equiv \int (a - A^a_{ss}) \psi_{ss} u'(c_{ss})\, dx$ *and* $\Psi^b_{ss} \equiv \int (b - A^b_{ss}) \psi_{ss} \partial_b W_{ss}\, dx$ *are the wealth- and bond-weighted commitment premiums. The promise multiplier* ψ_{ss} *solves the stationary forward equation.*

Commitment grafts the promise premiums Ψ^a_{ss} (inflation channel) and Ψ^b_{ss} (portfolio channel) onto the discretion conditions; the final term in (242) is the promise-weighted deposit-cost term, using $\partial_b W_{ss} - \partial_a W_{ss} = \Phi'(d_{ss}) u'(c_{ss})$ from (225). Setting $\psi_{ss} \equiv 0$ recovers Proposition 13.

G.3 Optimal separated policy with adjustment costs

In the separated problem of Section 4 the central bank chooses $(\dot{M}_s, \dot{B}^{CB}_s)$ with zero remittances, facing a Treasury that runs the passive rule $T_t = \bar{T} + \kappa(\bar{b} - \phi_t B^T_t)$ from (49); public bonds are $B = B^T - B^{CB}$, with real values $A^{b,T} = \phi B^T$ and $A^{b,CB} = \phi B^{CB}$. The consolidated budget splits into the Treasury budget (5), with shadow price ζ_s, and the central bank budget (6), with shadow price ϑ_s. The household block, wedges (235), and forward equation for ψ_s are unchanged; the social-value HJB prices transfers at the Treasury value $-\zeta_s T_s$, and ϑ_s is now the central bank's shadow price. The fiscal rule makes ζ_s its own co-state,

$$\dot{\zeta}_s = (\rho - r^b_s)\, \zeta_s - \varphi_s + \kappa(\zeta_s - \theta_s), \qquad \theta_s \equiv \int \partial_a \gamma_s\, dG_s + \int \psi_s\, \partial_a W_s\, dx, \tag{243}$$

where θ_s is the social value of a uniform transfer; the δg^M variation gives the separated revaluation $R_s = \vartheta_s (A^a_s - A^{b,CB}_s) + \zeta_s A^{b,T}_s$, and the bond-rate variation gives

$$\int b\, \partial_b \gamma_s\, dG_s + \int b\, \psi_s \partial_b W_s\, dx = \zeta_s A^{b,T}_s - \vartheta_s A^{b,CB}_s. \tag{244}$$

The bond-rate condition retains its household-drift term, the direct two-asset analog of the consolidated i^b FOC: the household benefit of a higher illiquid return is balanced against the Treasury's interest cost (valued at ζ_s) net of the central bank's interest income (valued at ϑ_s).

Discretion. Under full discretion ($\psi_s \equiv 0$), $\theta_{ss} = \int \partial_a \gamma_{ss}\, dG_{ss}$, and (243) in steady state gives the convex combination

$$\zeta_{ss} = \frac{(\rho - r^b_{ss})\, \vartheta_{ss} + \kappa\, \theta_{ss}}{(\rho - r^b_{ss}) + \kappa}, \tag{245}$$

so ζ_{ss} lies between the central bank's ϑ_{ss} and the household value θ_{ss}, and the wedges $\vartheta_{ss} - \theta_{ss}$, $\zeta_{ss} - \theta_{ss}$ share a sign.

Proposition 15 (Optimal separated policy under full discretion, adjustment costs)**.** *Under Assumptions 3 and 4 and full discretion with the passive fiscal rule (49), the optimal steady state is characterized by the liquid-wealth and illiquid-wealth margins*

$$\int a\, \partial_a \gamma_{ss}\, dG_{ss} = \vartheta_{ss} A^a_{ss}, \qquad \int b\, \partial_b \gamma_{ss}\, dG_{ss} = \zeta_{ss} A^{b,T}_{ss} - \vartheta_{ss} A^{b,CB}_{ss}, \tag{246}$$

with ζ_{ss} from (245). The illiquid margin is the bond-rate (i^b) condition; the inflation (δg^M) condition is their sum, $R_{ss} = \vartheta_{ss}(A^a_{ss} - A^{b,CB}_{ss}) + \zeta_{ss} A^{b,T}_{ss}$. Equivalently, centering at the uniform-transfer value θ_{ss},

$$\mathrm{Cov}_{G_{ss}}(a, \partial_a \gamma_{ss}) = (\vartheta_{ss} - \theta_{ss}) A^a_{ss}, \qquad \int b\, \partial_b \gamma_{ss}\, dG_{ss} - \theta_{ss} A^b_{ss} = (\zeta_{ss} - \theta_{ss}) A^{b,T}_{ss} - (\vartheta_{ss} - \theta_{ss}) A^{b,CB}_{ss}. \tag{247}$$

The two conditions are the separated analogs of the consolidated discretion conditions, with the consolidated ϑ replaced by the balance-sheet pair $(\vartheta_{ss}, \zeta_{ss})$: inflation revalues the central bank's net money issuance at ϑ_{ss} and the Treasury's diluted debt at ζ_{ss}, while the liquidity premium trades the Treasury's interest cost against the central bank's interest income. Summing the two left-hand sides of (247) reproduces the single-asset separated covariance $(\vartheta_{ss} - \theta_{ss})(A^a_{ss} - A^{b,CB}_{ss}) + (\zeta_{ss} - \theta_{ss}) A^{b,T}_{ss}$ of the main body, and setting $\vartheta_{ss} = \zeta_{ss} = \theta_{ss}$ collapses them to the consolidated conditions (239). The sign of each wedge tracks the fiscal stance through the transfer condition $\zeta_{ss} = \theta_{ss}$ that the rule forgoes.

Commitment. Under full commitment the promise multiplier reappears, $\theta_{ss} = \int \partial_a \gamma_{ss}\, dG_{ss} + \int \psi_{ss} \partial_a W_{ss}\, dx$, and the premiums Ψ^a_{ss}, Ψ^b_{ss} enter exactly as in the consolidated case.

Proposition 16 (Optimal separated policy under full commitment, adjustment costs)**.** *Under Assumptions 3 and 4 and full commitment with the passive fiscal rule (49), the stationary equilibrium is characterized by*

$$\mathrm{Cov}_{G_{ss}}(a, \partial_a \gamma_{ss}) + \Psi^a_{ss} = (\vartheta_{ss} - \theta_{ss}) A^a_{ss}, \tag{248}$$

$$\Big[\int b\, \partial_b \gamma_{ss}\, dG_{ss} + \int b\, \psi_{ss} \partial_b W_{ss}\, dx \Big] - \theta_{ss} A^b_{ss} = (\zeta_{ss} - \theta_{ss}) A^{b,T}_{ss} - (\vartheta_{ss} - \theta_{ss}) A^{b,CB}_{ss}, \tag{249}$$

with $\Psi^a_{ss} = \int (a - A^a_{ss}) \psi_{ss} u'(c_{ss})\, dx$ and the bond-weighted promise contribution on the left of (249), and ζ_{ss} from (243).

Commitment changes the separated problem exactly as it changes the consolidated one: it adds the wealth- and bond-weighted promise premiums to the two margins, leaving the balance-

sheet structure ($\vartheta_{ss} - \theta_{ss}, \zeta_{ss} - \theta_{ss}$) intact. The four cases thus differ only through two switches—consolidation ($\vartheta = \zeta = \theta$) and commitment ($\psi \equiv 0$)—applied to the same pair of liquid and illiquid optimality conditions.

H One-asset model: money only

This appendix solves the planner's problem when money is the only publicly supplied asset, $B_t \equiv 0$. With no bonds there is no liquidity premium and no portfolio choice: a household's real wealth $\omega \geq 0$ is its real money balance, and the only monetary instrument is the growth rate of money. The model is a strict simplification of the baseline, and the optimal-inflation condition is unchanged; the bond-rate condition has no counterpart. The monetary–fiscal separation of Section 4 also has no counterpart—with no debt the government budget balances period-by-period, so the separated and consolidated problems coincide—so we treat only the consolidated case.

Environment. A household with real wealth (real money) ω and productivity z solves

$$\rho W_t(\omega, z) = \max_{c_t, y_t} \; u(c_t) + \alpha U(y_t) + \mathcal{A}_t[W_t] + \partial_t W_t, \qquad c_t \geq 0, \; 0 \leq y_t \leq \omega, \tag{250}$$

with generator $\mathcal{A}_t[W] = [r_t^a \omega + z_t - c_t + \mathcal{T}_t]\partial_\omega W + \alpha[W(\omega - y_t, z) - W] + \mathcal{A}_z W$ and $r_t^a = \dot{\phi}_t/\phi_t$ the real return on money. The two household first-order conditions are

$$u'(c_t) = \partial_\omega W_t, \qquad U'(y_t) = \partial_\omega W_t(\omega - y_t, z), \tag{251}$$

the lumpy condition holding with equality on the interior and at the corner $y_t = \omega$ in the full-depletion region. Relative to the baseline, the lumpy condition carries no liquidity-premium term: with no bond to forgo, the only cost of the lumpy purchase is the marginal value of the wealth it consumes. The distribution G_t over (ω, z) solves the Kolmogorov forward equation $\partial_t G_t = \mathcal{A}_t^*[G_t]$, with a single market-clearing condition and the seigniorage-financed government budget

$$\int \omega \, dG_t = \phi_t M_t \equiv A_t, \qquad \int \mathcal{T}_t \, dG_t = \phi_t \dot{M}_t. \tag{252}$$

Optimality system. The optimality system is obtained discretize-first, following Appendix C with the bond, liquidity-premium, and bond-clearing terms deleted; the derivation is collected in a companion note available upon request. It introduces the social value $\gamma_s(\omega, z)$, the promise multiplier ψ_s, the shadow price of government revenue θ_s, and the consumption and lumpy wedges $\tilde{\chi}_s \equiv \partial_\omega \gamma_s - u'(c_s)$ and $\tilde{\eta}_s \equiv \partial_\omega \gamma_s(\omega - y_s, z) - U'(y_s)$ on the first-order conditions (251). Under lump-sum transfers the shadow price satisfies $\theta_s = \int \partial_\omega \gamma_s \, dG_s + \int \psi_s \partial_\omega W_s \, dx$, and the inflation margin is governed by the wealth-weighted social value $R_s \equiv \int \omega \, \partial_\omega \gamma_s \, dG_s + \int \omega \, \psi_s \partial_\omega W_s \, dx$ through the static revaluation $R_s = \theta_s A_s$. There is no bond-rate condition.

Proposition 17 (Optimal money-only policy under full discretion)**.** *Under full discretion with lump-sum transfers, the optimal steady state satisfies $r^{a*}_{ss} < \rho$ and the inflation condition*

$$\int \omega\, \partial_\omega \gamma_{ss}\, dG_{ss} = \theta_{ss} A_{ss}, \qquad \theta_{ss} = \int \partial_\omega \gamma_{ss}\, dG_{ss}, \tag{253}$$

equivalently, in covariance form,

$$\mathrm{Cov}_{G_{ss}}(\omega, \partial_\omega \gamma_{ss}) = 0. \tag{254}$$

This is exactly the consolidated baseline inflation condition (Proposition 3): the inflation tax is set so that wealth and its social marginal value are uncorrelated, trading the distortion of money demand against the redistribution the inflation tax achieves. The only change is that the baseline bond-rate condition $\mathrm{Cov}_{G_{ss}}(\omega - y, \partial_\omega \gamma_{ss}) = \int \eta_{ss} u'(c_{ss})\, dx$ has no counterpart, so the single instrument r^a_{ss} is pinned by (254) alone; the Friedman rule $r^a_{ss} = \rho$ obtains only when transfers can be made state-contingent.

Proposition 18 (Optimal money-only policy under full commitment)**.** *Under full commitment with lump-sum transfers, the stationary equilibrium satisfies*

$$\int [\omega\, \partial_\omega \gamma_{ss}\, g_{ss} + \omega\, \psi_{ss} \partial_\omega W_{ss}]\, dx = \theta_{ss} A_{ss}, \qquad \theta_{ss} = \int \partial_\omega \gamma_{ss}\, dG_{ss} + \Psi_{ss}, \tag{255}$$

with commitment premium $\Psi_{ss} \equiv \int \psi_{ss} u'(c_{ss})\, dx$; equivalently, in covariance form,

$$\mathrm{Cov}_{G_{ss}}(\omega, \partial_\omega \gamma_{ss}) + \Psi^\omega_{ss} = 0, \qquad \Psi^\omega_{ss} \equiv \int (\omega - A_{ss}) \psi_{ss} u'(c_{ss})\, dx. \tag{256}$$

The promise multiplier ψ_{ss} solves the stationary forward equation, which carries only the consumption and lumpy sources.

Commitment adds the wealth-weighted promise premium Ψ^ω_{ss} to the discretion condition, exactly as in the consolidated baseline (Proposition 4); setting $\psi_{ss} \equiv 0$ recovers Proposition 17. Dropping bonds thus removes the entire bond-rate margin and leaves the inflation margin untouched.