The long and short of financing government spending*

Jochen Mankart†, Rigas Oikonomou‡ and Romanos Priftis§

June 15, 2026

Abstract

We study optimal policy in a heterogeneous-agent New Keynesian model with short- and long-term government debt, where short-term bonds provide liquidity and safety services to the private sector. By satisfying households' demand for precautionary savings, short-term bonds command a liquidity/safety premium and therefore earn a lower equilibrium return. Households do not require full compensation for inflation to hold this debt, and its real return deviates from consumption growth. We show that the Ramsey-optimal policy features a positive, time-varying inflation target. The government also finds it optimal to introduce volatility in labor taxes to reduce the real return on newly issued short-term debt through novel intertemporal tax-distortion channels. In addition, short-term debt plays an active role in fiscal stabilization: the government adjusts its supply over the business cycle, using it as a shock absorber alongside long-term debt rather than merely as a funding vehicle for long-bond positions. Despite its insurance benefits, however, short-term debt is never supplied in quantities sufficient to eliminate financial frictions. Finally, we show that the Ramsey problem admits multiple locally optimal allocations. We develop a computational strategy that robustly handles this multiplicity and delivers a complete characterization of optimal policy.

Keywords: Monetary Policy, Fiscal Policy, Debt Maturity, Liquidity and Safety of Debt, Financial Frictions, Computational Methods.
JEL classifications: C63, E43, E62, H63

*We received helpful comments from, Arpad Abraham, Fernando Broner, Harris Dellas, Refet Gürkaynak, Tasos Karantounias, Olivier Loissel, Mathias Klein, Jean-Baptiste Michau, Kurt Mitman, Ricardo Nunes, Gert Peersman, Ansgar Rannenberg, Kjetil Storesletten, Harald Uhlig, Dimitri Vayanos, Gianluca Violante, Raf Wouters and seminar participants at the Deutsche Bundesbank, the National Bank of Belgium, the Bank of Spain, at Bristol, Surrey, AUEB, UC3M, and Crest/Ecole Polytechnique. The views expressed in this document are exclusively the opinion of the authors and do not necessarily represent those of the Bank of Spain, the Bundesbank, the European Central Bank, the National Bank of Slovakia or any other official institution. Oikonomou is grateful to the National Bank of Belgium for funding.

†National Bank of Slovakia & Deutsche Bundesbank; jochen.mankart@nbs.sk
‡UC Louvain & University of Surrey; rigas.oikonomou@uclouvain.be
§European Central Bank; romanos.priftis@ecb.europa.eu

1 Introduction

Understanding how governments can optimally finance public expenditures is central to macroeconomic policy. A large literature addresses this question within the canonical Ramsey framework, which features a fully committed planner, stochastic government spending, and distortionary taxes and inflation. This framework yields sharp prescriptions. Tax rates should be smoothed over time, and follow a near-random walk after fiscal shocks (Barro, 1979; Aiyagari, Marcet, Sargent, and Seppälä, 2002). Inflation should deviate only temporarily from its zero steady-state target and vary less than taxes, due to its higher distortionary costs (Schmitt-Grohé and Uribe, 2004a; Faraglia, Marcet, Oikonomou, and Scott, 2013; Leeper and Zhou, 2021). Debt should be issued in long maturities because (i) long bond prices drop in high-spending states and thus can absorb fiscal shocks (Angeletos, 2002; Buera and Nicolini, 2004), and (ii) a lengthier maturity allows the government to spread inflation adjustments over a longer period, thereby achieving a smoother path (Lustig, Sleet, and Yeltekin, 2008; Leeper and Zhou, 2021).

While influential, this benchmark abstracts from an increasingly emphasized role of government debt in modern economies: its function as a source of safe and liquid assets for the private sector. A growing literature shows that government bonds are valued not only for their returns but also for the liquidity and safety services they provide (Angeletos, Collard, and Dellas, 2022; Hagedorn, 2018; Mian, Straub, and Sufi, 2025; Auclert, Cai, Rognlie, and Straub, 2024, among others). In environments with financial frictions/incomplete markets, households use these assets for self-insurance and consumption smoothing. A higher supply of debt improves insurance against idiosyncratic risk and mitigates the impact of frictions.

We depart from the canonical Ramsey framework by allowing government debt to provide liquidity and safety, assuming that (due to these attributes) households use government bonds to self-insure against idiosyncratic consumption risk. We emphasize the role of short-term bonds as providers of safety and liquidity, while maintaining that longer maturities, as in the benchmark Ramsey model, are valued primarily for their return properties, facilitating the intertemporal allocation of consumption (Greenwood, Hanson, and Stein, 2015). Although our analysis can be extended to alternative definitions of short- and long-term maturities, we adopt the simplifying assumption—made for tractability given the complexity of the policy problem we solve—that short-term debt is one-period debt and long-term bonds are perpetuities.

We show that, in this setting, the classic Ramsey prescriptions for inflation, taxation, and debt management are significantly altered. First, optimal policy features a strictly *positive and time-varying inflation target*: following adverse fiscal shocks, the planner optimally commits to a permanently higher inflation rate. Second, the optimal maturity structure is not tilted exclusively toward long-term debt. Instead, the government optimally supplies a substantial amount of short-term debt because it provides liquidity and safety services. At the same time, the planner never supplies this debt in quantities sufficient to (effectively) eliminate financial frictions. Finally, the government

finds it optimal to introduce volatility in labor taxes to depress the real return on newly issued short-term debt, operating through novel channels of intertemporal tax distortions which do not have a counterpart in the canonical Ramsey framework.

In Section 2, we describe our baseline model, which can be viewed as an extension of Hagedorn (2018). As in that paper, we use a simple microfoundation for liquidity and safety services, assuming that households face idiosyncratic shocks to the marginal utility of consumption, making them ex post heterogeneous in their spending needs. Households use short-term bonds to finance spending. Moreover, they pool resources at the end of each period, which makes heterogeneity temporary. Our setting follows a recent strand of the literature developing tractable models to analyze the effect of precautionary savings (see Hagedorn, 2018; Mian et al., 2025; Angeletos et al., 2022).

Even in this simple environment, characterizing optimal policy is challenging: the program features multiple local optima, with several allocations satisfying the planner's first-order conditions. In particular, multiplicity arises because these conditions are satisfied in two distinct regions. First, in a region with *binding financial frictions*, where the supply of short debt is limited and households are constrained. Second, when the supply of debt is sufficiently abundant, the model becomes isomorphic to the canonical Ramsey framework, in which government bonds do not provide liquidity/safety. Along the transition to the long-run equilibrium, or with recurrent fiscal shocks, the planner's solution may fluctuate across regimes, with optimal policy alternating between states of binding frictions and states of liquid/safe asset abundance. In this environment, obtaining robust and economically meaningful characterizations of optimal policy is difficult.

We first derive an analytical result based on the date-0 allocation in a simplified environment with quasilinear preferences and rigid prices—a setting commonly used in the literature (e.g., Aiyagari et al., 2002; Angeletos et al., 2022). We show that, away from the first best, the optimal policy remains in the region of binding financial frictions. By limiting the supply of short-term debt, the planner extracts liquidity/safety rents and reduces distortionary taxation (Angeletos et al., 2022).

Beyond this simple setting, when preferences are not quasilinear and prices are not fully rigid, analytical solutions are generally infeasible. In such environments, the economy exhibits transitional dynamics, as the planner distorts taxes and inflation intertemporally to manage the debt burden. Solving the model using global methods, we find that oscillatory transitions across regimes are still not optimal.

Building on these findings, we study the properties of optimal inflation and taxation in the equilibrium with financial frictions (Section 3). We first show that the optimal policy features a positive, time-varying inflation target. When debt provides liquidity/safety, agents do not require full compensation for expected inflation to hold it. The planner exploits this by setting a positive average (steady-state) inflation rate and committing to adjust the target upward (downward) following positive (negative) fiscal shocks.

These properties stand in sharp contrast to the canonical Ramsey model, in which deviations of inflation from its zero steady-state target are only temporary and fully determined by the matu-

rity structure. In that setting, inflation is pinned down by the government's intertemporal budget constraint and adjusts to equate the market value of debt with the present value of fiscal surpluses. Permanent deviations of inflation from target are not optimal and are in fact wasteful from the perspective of debt solvency (unless debt is issued as consols Leeper and Zhou, 2021).

With financial frictions, this logic breaks down. Part of government revenue takes the form of rents from liquidity/safety provision, and inflation appears on both sides of the intertemporal debt constraint, affecting both the real value of debt and the present value of surpluses. The intertemporal constraint no longer fully determines inflation (Hagedorn, 2018, 2024). Following a spending shock, the planner raises the inflation target to increase rents, even when debt is only short-term.

We then turn to the implications of our model for optimal taxation. Financial frictions create novel intertemporal tax distortions. As in the canonical framework, the planner manipulates the timing of taxes to alter real discount rates and reduce the cost of debt issuance. However, whereas in the canonical model these distortions are purely temporary—operating through short-lived movements in short bond prices —this is not the case with financial frictions. Just like with inflation, persistence obtains because bonds are valued for liquidity/safety and households are willing to hold them despite low real returns. The planner responds by introducing higher intertemporal variability in taxes. Furthermore, financial frictions create an additional channel through which tax fluctuations affect the economy. Changes in taxes alter the fractions of constrained and unconstrained households, allowing the planner to influence the relative returns these groups derive from holding government debt. This generates a novel source of intertemporal tax volatility that has no counterpart in the canonical model.

Using quantitative simulations, we explore the importance of these channels. Although they can have significant effects on inflation and tax outcomes in some parameterizations, the prediction that taxes are substantially more volatile than inflation remains robust.

In Section 4 we study the optimal composition of public debt. We begin by outlining our numerical approach, which addresses two significant challenges: the high-dimensional state space and the inherent multiplicity of equilibrium solutions.

Specifically, since long-term debt takes the form of coupon payments extending over an infinite horizon, the state vector is, in principle, infinite-dimensional. The planner's commitments extend indefinitely into the future; consequently, an infinite number of lags of Lagrange multipliers, bond holdings, and inflation are required as state variables. Assuming that long bonds are perpetuities with geometrically decaying coupons allows us to collapse this infinite-dimensional structure into a single recursive state variable. Furthermore, because liquid/safe debt is one-period, we do not need to track multiple lags of the state variable capturing financial conditions, the ratio of debt to consumption in the model. Despite these simplifications, the model still features a high-dimensional state space, rendering the solution of the saddle-point functional equation (Marcet and Marimon, 2019) on a discretized grid computationally intractable.

We solve the model efficiently using the planner's first-order conditions, applying the Parameter-

ized expectations algorithm (PEA, Den Haan and Marcet, 1990). Doing so, however, means that we need to confront the second difficulty, the multiplicity of local optima. We augment our algorithm with an approximation of the planner's welfare objective function to be able to discern the global optimum among the multiple critical points that satisfy the first-order conditions. This approach of merging simulations and value function evaluation is similar to Judd, Maliar, and Maliar (2015) but has so far not been applied to a Ramsey problem with non-convexities.

Applying our algorithm, we investigate how the standard prescription for the maturity structure of debt changes when short- and long-term securities serve fundamentally different roles. In particular, we examine how the planner trades off the rents from restricting the supply of liquid/safe assets against the welfare costs associated with imperfect household insurance; and whether some spending histories make it optimal to supply these assets abundantly eliminating financial frictions.

In our baseline specification, government debt is constrained to be non-negative (e.g., Lustig et al., 2008; Nosbusch, 2008; Faraglia, Marcet, Oikonomou, and Scott, 2019, among others). Within this benchmark, the optimal policy never entails supplying liquid/safe assets in quantities sufficient to eliminate frictions, in line with our analytical result.

In a stochastic economy, the government has even weaker incentives to target a high supply of short-term bonds. Intuitively, since long bond prices tend to decline in high-spending states, long-debt provides a hedging benefit to the government, allowing to smooth taxes and inflation across states. This introduces a trade-off between issuing long-term debt, which helps insure the government against aggregate fiscal shocks, and issuing short-term debt, to provide households with a valuable instrument for self-insurance against idiosyncratic risk. This trade-off gives an additional reason, beyond rent extraction, to limit the supply of liquid safe assets.

Furthermore, the government actively manages the supply of short-term bonds over the business cycle. Following an increase in spending, the optimal response is to reduce short-term debt and rely more heavily on long-term issuance. Through this channel, the planner increases the value of rents extracted and, effectively, turns short-term bonds into a state-contingent claim, in a manner analogous to the price-driven hedging of long debt. This mechanism, which overturns the canonical model's property that short-term debt mainly serves as a passive funding vehicle for long debt positions, implies a rising share of long-term debt during periods of rapidly increasing government indebtedness—a pattern consistent with the U.S. data.

Finally, we extend our analysis to allow the government to invest in private assets (i.e., to take negative debt positions). Provided empirically relevant limits to the size of positions the government can take in private asset markets, our findings continue to hold.

Related Literature. Our paper sits at the intersection of two strands of literature. The first studies optimal policy in environments where government debt is net wealth and provides safety and liquidity services to households, satisfying their demand for precautionary savings (Angeletos et al., 2022; Auclert et al., 2024; Bhandari, Evans, Golosov, Sargent, et al., 2019; Dávila and Schaab, 2023; Le Grand, Martin-Baillon, and Ragot, 2026; Acikgöz, Hagedorn, Holter, and Wang, 2018, among

others).

As discussed, we follow this literature in modeling safety and liquidity within a simple framework. We abstract from heterogeneity-driven considerations in optimal policy design. Importantly, our focus—the optimal composition of government debt—differs from that of the existing literature, motivating the use of a parsimonious model.[1]

An important precursor to our analysis is Angeletos et al. (2022), who characterize optimal Ramsey tax and debt policy in an economy where all government debt is safe and liquid. As in our model, the optimal policy must be characterized across two distinct regions: one in which financial frictions bind and one in which they do not. A key result is that, near the equilibrium with binding frictions, the economy exhibits transitional dynamics with gradual convergence to the steady state. Along this path, optimal taxes differ markedly from the canonical Ramsey benchmark.

Our model predicts different outcomes. First, long-term debt eliminates the transitional dynamics emphasized by Angeletos et al. (2022). Because the government can freely adjust long debt, it can instantly target any desired level of liquid/safe debt supply. This restores classical tax smoothing in the region of binding financial frictions. However, it also introduces an additional complication: optimal policy may oscillate between liquidity abundance and scarcity. Addressing this possibility is one of the central challenges of our paper.[2]

Finally, we study the optimal policy equilibrium in a stochastic economy with sticky prices and curved utility, extending Angeletos et al. (2022). Whereas their main focus is on how debt supply affects rents and the level of tax distortions, we primarily consider an economy in which the supply of liquid/safe short debt can be used to smooth taxes and inflation across states, together with the standard state-contingent payoffs of long bonds.

The second strand of literature studies the optimal composition of public debt in macroeconomic models (see, e.g., Angeletos, 2002; Buera and Nicolini, 2004; Nosbusch, 2008; Lustig et al., 2008; Faraglia, Marcet, and Scott, 2010; Faraglia et al., 2019; Greenwood et al., 2015; Debortoli, Nunes, and Yared, 2017; Aparisi de Lannoy et al., 2025; Bigio, Nuño, and Passadore, 2023; Chafwehé, de Beauffort, and Oikonomou, 2026). Angeletos (2002) and Buera and Nicolini (2004) show that, to achieve tax smoothing, an optimizing government issues long-term debt while holding short-term assets, exploiting the negative covariance between fiscal deficits and long bond prices. Lustig et al. (2008) and Nosbusch (2008) abstract from government investment in private assets and show that

[1]Most papers studying optimal policy with heterogeneity rely on perturbation-based methods (e.g., Auclert et al., 2024; Bhandari et al., 2019; Dávila and Schaab, 2023; Aparisi de Lannoy, Bhandari, Evans, Golosov, and Sargent, 2025). Such methods are not well suited to addressing the multiplicity of equilibria inherent in our model. Moreover, at least when restricted to first- or second-order approximations, they do not naturally accommodate the analysis of time-varying portfolios in environments where public debt follows a random walk and may drift far from the point of approximation.

[2]See also Canzoneri, Cumby, and Diba (2016) and Sims (2026). Canzoneri et al. (2016) use a model with liquid government debt and money. Assuming complete financial markets and flexible prices, they prove that the optimal policy over the business cycle is to provide abundant liquidity in both debt and money. We reach the opposite conclusion in a model with non-contingent debt and sticky prices. Sims (2026) argues that the optimality of abundant liquidity provision hinges on the government's ability to invest in private assets. This is consistent with our analysis in the final subsection of the paper. Unlike Sims (2026), we study optimal policy in a stochastic economy.

even in this restricted setting, it is optimal to rely predominantly on long-term debt.

Recent work introduces frictions that rationalize short-term debt issuance. For example, Debortoli et al. (2017) emphasize lack of commitment and show that a flat maturity structure mitigates time inconsistency. Faraglia et al. (2019) show that costly long-term debt repurchases make short-term debt useful for tax smoothing. Aparisi de Lannoy et al. (2025) show that short-term debt is desirable when government budget risk is driven by fluctuations in real interest rates.

Our paper complements this work. In our model, the optimizing government always issues some short-term debt because it provides liquidity/safety to households. This mechanism is related to Greenwood et al. (2015), who develop a stylized three-period model with liquidity preferences and study how debt maturity can hedge exogenous interest-rate risk. We instead analyze optimal policy in a fully fledged New Keynesian framework with distortionary taxes and sticky prices, focusing on spending shocks. Calibrating our model to match the empirical evidence of Greenwood et al. (2015), we capture the distinct roles of short- and long-term government bonds in a simplified way. As in the above papers, we compare our results to U.S. data and show that a model in which short-term debt provides hedging benefits to the private sector can account for several salient features of public debt portfolios.

2 Baseline Model and Optimal Policy

This section presents the model and the optimal policy problem. We first introduce the baseline environment in which short-term debt provides liquidity and safety services. We then set up the Ramsey policy program and derive the optimality conditions.

2.1 The baseline model

Our baseline model extends the one-asset economy of Hagedorn (2018). Government debt provides liquidity/safety services by allowing households to finance unexpected spending needs and smooth consumption. Our extension is the introduction of two assets—short-term and long-term government bonds. We assume that only short-term debt provides liquidity/safety services. Before presenting the model equations, we briefly motivate this assumption and refer the reader to Ghomi, Mankart, Priftis, and Oikonomou (2026) for a more extensive discussion.

Our modeling choice is grounded in a broad literature documenting that short-term government bonds are safer assets and provide money-like services to the private sector.

First, the literature emphasizes that long-term government bonds are riskier stores of value than short-term bonds. Their valuations are more sensitive to fluctuations in expected inflation and discount rates. Standard asset-pricing models imply that long-duration bonds command higher expected returns precisely because they are exposed to substantial interim repricing risks (see, for example, Rudebusch and Swanson, 2012; Piazzesi, Schneider, Benigno, and Campbell, 2007; Aparisi de Lan-

noy et al., 2025). Although these risks are not modeled explicitly in our framework, they provide a rationale for why households facing idiosyncratic risk would prefer to accumulate precautionary savings in short-duration assets. When agents are forced to liquidate positions following adverse shocks, short-term bonds preserve value more reliably and therefore provide better insurance against unexpected consumption needs.[3]

Second, many papers argue that short-term government debt provides liquidity to the private sector, relative to longer-maturity bonds. Particularly relevant to our paper, Greenwood et al. (2015) provide empirical evidence that short-term government debt functions as money. We return to this evidence in the calibration section. On the theoretical side, numerous papers microfound the superior liquidity of short-term government debt, emphasizing its role in facilitating transactions or serving as collateral. See, for example, Bansal and Coleman (1996), Geromichalos, Lee, and Oikawa (2016), Williamson (2016), and Bigio et al. (2023).[4] The notion that more liquid assets are better suited for precautionary savings is both intuitive and widespread in the literature.

With our reduced-form approach, we aim to jointly capture the safety and liquidity attributes of short-term government debt without committing to a specific underlying microfoundation. Moreover, these attributes are closely intertwined (e.g., Geromichalos, Herrenbrueck, and Lee, 2023), making a sharp distinction between them less compelling, at least in the present context.

We emphasize that, although our model distinguishes between one-period short-term debt and perpetuities for tractability, our results do not rely on this specific maturity structure. What matters for our mechanism is the distinction between government liabilities that provide liquidity and safety services and those that primarily facilitate intertemporal transfers of consumption. Our results would extend to an environment in which both short- and medium-term government debt provide liquidity and safety, while sufficiently long-duration debt remains exposed to repricing risk and therefore

[3] Challe, Le Grand, and Ragot (2013) use a related argument to formalize the notion of a "liquidation risk premium": households hit by adverse idiosyncratic income shocks may need to liquidate their portfolios precisely when bond prices are depressed. See also Schneider (2022) for a model in which risk sharing among heterogeneous households is central to bond pricing.

[4] In particular, Bansal and Coleman (1996) develop a model in which safe and liquid Treasury bills are used by banks to back checkable deposits and use the model to explain the positive slope of the yield curve. See also Gorton and Metrick (2012) for a related perspective in which short debt backs privately issued money-like claims.

Geromichalos et al. (2016) and Williamson (2016) provide microfoundations based on search-theoretic frameworks. In Geromichalos et al. (2016), short-term bonds are more useful for financing unexpected consumption needs because "off-the-run" long-term bonds can only be liquidated in secondary markets subject to search frictions. Williamson (2016) instead emphasizes collateral differences, assuming that short-term debt is more pledgeable than long-term debt. Finally, Bigio et al. (2023) study a model with financial intermediaries and investors in which long-term bond markets are less liquid because of search frictions.

Some of these mechanisms are closely related to our modeling assumptions, particularly the framework of Geromichalos et al. (2016). Moreover, if, as in Bansal and Coleman (1996), short-term debt influences the supply of credit, then it directly affects households' ability to hedge against idiosyncratic shocks. More generally, for theories emphasizing the collateral and transaction services of short-term debt, our framework can be interpreted as a reduced-form representation of these liquidity services, in the spirit of Greenwood et al. (2015), Angeletos et al. (2022) and Sims (2026).

Finally, long-term debt may be less liquid because it is often held for retirement savings. Such assets are frequently held in accounts with withdrawal or rebalancing costs. Although our model abstracts from life-cycle considerations. This interpretation is consistent with our assumption that long-term bonds primarily transfer consumption across periods rather than satisfy short-run liquidity needs.

provides these services less effectively.

2.1.1 Timing and preferences

We provide a brief description of the model, focusing on key equations, and relegate derivations to the appendix. A more detailed discussion of the frictions and the equilibrium can be found in Hagedorn (2018).

The economy is populated by a continuum of infinitely lived, ex-ante identical agents/households. Time is discrete, and each period t is divided in two subperiods, t_1, t_2.[5]

The timing of events is as follows. In subperiod 1, households choose consumption C_t, savings, and labor supply. In subperiod 2, each household chooses consumption c_t after experiencing an idiosyncratic preference shock that makes her desired consumption differ from that of other households. We assume that consumption in t_2 can only be financed by short-term assets chosen in t_1. To preserve tractability, we follow Hagedorn (2018) and assume that households belong to large families that pool resources and redistribute through transfers at the end of subperiod 2. Consequently, all agents begin each period with the same wealth level and make identical consumption and portfolio decisions.

The preferences of household i (when the shocks have been revealed) are:

$$u(C_t^i) + \theta^i v(c_t^i) - \chi \frac{{h_t}^{i,1+\gamma}}{1+\gamma}$$

$\theta^i \in [0, \overline{\theta}]$ is the household-specific preference shock that scales utility derived from consumption in t_2. A high realization of θ^i implies high expenditure needs and therefore a high consumption demand. We assume θ^i is i.i.d. across households and time, with density f and cumulative distribution function F.[6] We will consider cases where the distribution F has a bounded support ($\overline{\theta}$ is finite) and cases where the support of F is unbounded above ($\overline{\theta} = \infty$.)

h_t^i denotes hours worked by household i. Parameter χ affects the disutility of working and γ is the inverse of the Frisch elasticity of labor supply.

2.1.2 Assets and asset demand

In subperiod t_1, each household solves a portfolio choice problem, selecting optimal quantities of a short-term (S) and a long-term (L) bond. We denote by $B_{S,t}^i, B_{L,t}^i$ the nominal holdings of the two assets and let $b_{S,t}^i, b_{L,t}^i$ denote the real quantities (deflated by the price level P_t). Long bonds are

[5] Technically, t_1 and t_2 need not represent different points in real time; they are simply used to introduce the idea that households can participate in asset markets and make savings decisions (in t_1) before the full vector of state variables has been revealed. In our notation below, we very frequently condense t_1 and t_2 into t. We distinguish between t_1 and t_2 whenever it is absolutely necessary.

[6] Note that assuming that shocks are i.i.d. is necessary to rule out heterogeneity in portfolio choice decisions, when, say, agents experiencing a high θ^i today will likely expect a high θ^i tomorrow and have a stronger demand for short-term assets. For simplicity, we drop the superscript i for the θ shock from now on.

perpetuities paying geometrically decaying coupons. δ denotes the decay factor. The price of an L bond in period t is $q_{L,t}$.

Households purchase short-term bonds for two reasons: for their return (the inverse of the price $q_{S,t}$), and to finance consumption in subperiod 2. We assume that these consumption expenditures c_t^i are subject to the constraint $c_t^i \leq b_{S,t}^i$. Therefore, households desiring high expenditures in t_2 may be constrained by their short-term bond holdings chosen in t_1.

Crucially, in subperiod 2, each household can only use its own portfolio to finance c_t^i.[7] However, since households belong to families that pool resources after transactions are completed, they have identical wealth levels at the portfolio choice stage in t_1 and therefore hold identical quantities of short-term and long-term assets.

2.1.3 Household's problem

We now describe the household's program and the optimality conditions. The budget constraint in t_1 is:

$$(1) \quad P_t C_t^i + q_{L,t} B_{L,t}^i + q_{S,t} B_{S,t}^i = P_t(1-\tau_t) w_t h_t^i + (1 + q_{L,t}\delta) B_{L,t-1}^i + B_{S,t-1,2}^i + Div_t P_t - P_t \bar{C}_t^i$$

The left-hand side contains the household's choice variables. The first term on the right-hand side represents net labor income $(1-\tau_t) w_t h_t^i$ where w is the real wage rate and τ_t is the proportional income tax rate.

The terms $(1 + q_{L,t}\delta) B_{L,t-1}^i + B_{S,t-1,2}^i$ represent the nominal pay out from long- and short-term assets purchased in the previous period. Notice that $B_{S,t-1,2}^i$ has a subscript '2' used to denote that these are short bonds that remained in the household's portfolio after the transactions in subperiod 2 of period $t-1$ had been realized. The variable Div_t represents dividend income.[8]

$\bar{C}_t^i$ denotes the goods the household expects to sell to other families in subperiod 2. Importantly, $\bar{C}_t^i$ is not a choice variable for the household and it is used to ensure goods market clearing.[9] It holds that $E_\theta(c_t^i(\theta)) = \bar{C}_t^i$. The household will enter the next period with short-term bonds equal to

$$(2) \quad B_{S,t,2}^i = E_\theta(B_{S,t}^i - P_t(c_t^i(\theta)) + P_t \bar{C}_t^i,$$

[7] As explained in Hagedorn (2018) the interpretation of the uninsurability of the expenditure shock, θ, could then be a spatial one. In subperiod 2, family members are spatially separated and so the goods c_t^i have to be obtained from the market (from other families) in exchange for the asset.

[8] Firms earn profits and distribute dividends to households. We assume households hold identical equity shares across all firms and therefore receive identical dividends.

Moreover, we assume (as many papers do) that shares cannot be traded. This assumption is not restrictive since households are identical at the beginning of each period and would choose identical portfolios if trading were permitted.

Like long-term bonds, stocks cannot be used to finance subperiod 2 consumption, reflecting their lower liquidity and safety relative to short-term bonds.

[9] More specifically, since subperiods t_1 and t_2 may not represent different points in real time, households cannot distinguish between customers in t_1 and t_2 and how much is sold in either subperiod is exogenous. For details see Hagedorn (2018).

The household maximizes utility subject to (1), (2) and the constraint $c_t^i \leq b_{S,t}^i$. The Bellman equation is stated formally in the online appendix. We show that the following conditions hold at the optimum: First,

$$u'(C_t^i) = \theta v'(c_t^i) \quad \text{if} \quad \theta < \widetilde{\theta}_t \tag{3a}$$

$$c_t^i = b_{S,t}^i \quad \text{if} \quad \theta \geq \widetilde{\theta}_t \tag{3b}$$

When the realized value of θ falls below the threshold $\widetilde{\theta}_t$, the household's choice is unconstrained and satisfies $\theta v'(c_t^i) = u'(C_t^i)$. In contrast, if $\theta \geq \widetilde{\theta}_t$ the constraint binds and $c_t^i = b_{S,t}^i$. At the threshold, we have $\widetilde{\theta}_t v'(b_{S,t}^i) = u'(C_t^i)$.

Second, the optimal choice of short-term bonds yields:

$$q_{t,S} u'(C_t^i) = F_{\widetilde{\theta}_t} \beta E_t \frac{u'(C_{t+1}^i)}{\pi_{t+1}} + \int_{\widetilde{\theta}_t}^{\overline{\theta}} \theta v'(b_{S,t}^i) dF_\theta, \tag{4}$$

when $\widetilde{\theta}_t < \overline{\theta}$ and $q_{t,S} u'(C_t^i) = \beta E_t \frac{u'(C_{t+1}^i)}{\pi_{t+1}}$ when $\widetilde{\theta}_t \geq \overline{\theta}$.

Equation (4) equates the marginal utility cost of saving in short-term bonds, $q_{t,S} u'(C_t^i)$, with the marginal benefit of acquiring more of the asset. With probability $F_{\widetilde{\theta}_t}$, the shock is below the threshold value, and some short-term bonds will be carried over to the next period, yielding a real payoff $\frac{1}{\pi_{t+1}}$. If the shock exceeds the threshold, the constraint binds and all short-term bonds are sold to finance subperiod 2 consumption, a benefit captured by $\int_{\widetilde{\theta}_t}^{\infty} \theta v'(b_{S,t}^i) dF_\theta$.

Third, the price of the long-term bond satisfies a standard Euler equation:

$$q_{t,L} u'(C_t^i) = \beta E_t \frac{u'(C_{t+1}^i)}{\pi_{t+1}} (1 + \delta q_{t+1,L}) \tag{5}$$

Finally, the optimal choice of hours is determined by the familiar labor supply condition, $\chi \frac{(h_t^i)^\gamma}{u'(C_t^i)} = w_t(1 - \tau_t)$. For the remainder of this section, we suppress the index i and focus on an equilibrium in which households are ex ante identical in every period (Hagedorn, 2018).

2.1.4 Production / Government / Resource Constraints

We describe the production side of the model and the government.

Production. A final good is produced by aggregating infinitely many differentiated products. Each product is produced under monopolistic competition by a single producer operating a technology which is linear in the labor, $Y_t(j) = H_t(j)$.

The final good is given by $Y_t = \left(\int_0^1 Y_t(j)^{\frac{1+\eta}{\eta}} dj \right)^{\frac{\eta}{\eta-1}}$ where $\eta < -1$ governs the elasticity of substitution across the differentiated goods.

Producers maximize discounted profits subject to demand, taking labor costs w_t as given. Price adjustment involves resource costs à la Rotemberg (1982). In particular, $\text{Cost}_t = \frac{\omega}{2} \left(\frac{P_{jt}}{P_{jt-1}} - 1 \right)^2$, is

the cost that the firm has to bear whenever it changes its price. Parameter ω governs price rigidity; higher values imply steeper adjustment costs. When $\omega = 0$, prices are perfectly flexible.

In this standard setting, a symmetric equilibrium (all firms charge identical prices P_t and hire identical labor h_t) can be found (see, for example, Schmitt-Grohé and Uribe, 2004b). The model admits the New Keynesian Phillips curve

$$\pi_t(\pi_t - 1) = \frac{\eta}{\omega}\left(\frac{1+\eta}{\eta} - w_t\right)h_t + \beta E_t \frac{u'(C_{t+1})}{u'(C_t)}\pi_{t+1}(\pi_{t+1} - 1)$$

Government. The government levies taxes and issues short- and long-term debt to finance exogenous spending G_t. G_t is random and the only source of aggregate risk in the model. The government's nominal budget constraint is:

$$q_{t,S}B^g_{t,S} + q_{t,L}B^g_{t,L} = B^g_{t-1,S} + B^g_{t-1,L}(1 + \delta q_{t,L}) + P_t(G_t - w_t\tau_t h_t)$$

where the superscript g denotes government bond supply. Using market clearing to equate bond demand and supply, the government budget in real terms becomes:

$$q_{t,S}b_{S,t} + q_{t,L}b_{L,t} = \frac{b_{t-1,S}}{\pi_t} + \frac{b_{t-1,L}}{\pi_t}(1 + \delta q_{t,L}) + G_t - w_t\tau_t h_t$$

Resource Constraint. Combining the household and the government budget constraints yields the economy-wide resource constraint $C_t + \int c_t(\theta)dF_\theta + G_t + \frac{\omega}{2}(\pi_t - 1)^2 = Y_t$.

2.2 Optimal Policy

In this section, we set up the policy problem. Our approach closely follows a large body of literature studying optimal fiscal, monetary, and debt policies under a Ramsey planner (e.g., Aiyagari et al. (2002); Schmitt-Grohé and Uribe (2004b); Lustig et al. (2008); Faraglia, Marcet, Oikonomou, and Scott (2016); Faraglia et al. (2019); Angeletos et al. (2022), among others). As in these papers, we assume that a benevolent planner maximizes household utility subject to a set of sufficient implementability conditions for a competitive equilibrium.

For brevity, the formal definition of the equilibrium is provided in Appendix C, where we characterize the sequences of aggregate and household-specific variables that satisfy household and firm optimization, government policies, and market-clearing.

For the remainder of the paper, we focus on two utility specifications. First, in our baseline model, we set $u(\cdot) = v(\cdot) = \log(\cdot)$. Second, for some analytical derivations, we consider quasilinear preferences, with $u(C_t) = C_t$ and $v(\cdot) = \log(\cdot)$. We derive the Ramsey program under the baseline specification and defer the analogous formulation for the quasilinear case to appendix A. For notational convenience, we retain u' and u'' to denote the first and second derivatives of the utility function, respectively.

2.2.1 Policy objective and constraints

Following standard practice in Ramsey policy programs, we focus on a set of sufficient implementability conditions rather than the full set of equilibrium equations. The following proposition states the planner's constraints and derives the policy objective function.

Proposition 1 (Ramsey program) *The Ramsey planner solves the following program:*

$$\max_{\{\pi_t, Y_t, \tau_t, b_{L,t}, b_{S,t}, \widetilde{\theta}_t, C_t\}} E_0 \sum_{t\geq 0} \beta^t \widetilde{U}_t \tag{6}$$

$$\textit{where} \quad \widetilde{U}_t \equiv \left(1 + \mathbf{1}_{\widetilde{\theta}_t<\overline{\theta}} \int_0^{\widetilde{\theta}_t} \theta dF_\theta + \mathbf{1}_{\widetilde{\theta}_t\geq\overline{\theta}} \int_0^{\overline{\theta}} \theta dF_\theta\right) log(C_t)$$
$$+\mathbf{1}_{\widetilde{\theta}_t<\overline{\theta}}\left(\int_0^{\widetilde{\theta}_t} \theta log(\theta) dF_\theta + \int_{\widetilde{\theta}_t}^{\overline{\theta}} \theta log(b_{S,t}) dF_\theta\right) + \mathbf{1}_{\widetilde{\theta}_t\geq\overline{\theta}} \int_0^{\overline{\theta}} \theta log(\theta) dF_\theta - \chi \frac{Y_t^{1+\gamma}}{1+\gamma}$$

subject to constraints:

$$C_t\left(1 + \mathbf{1}_{\widetilde{\theta}_t<\overline{\theta}} \int_0^{\widetilde{\theta}_t} \theta dF_\theta + \mathbf{1}_{\widetilde{\theta}_t\geq\overline{\theta}} \int_0^{\overline{\theta}} \theta dF_\theta\right) + \mathbf{1}_{\widetilde{\theta}_t<\overline{\theta}} \int_{\widetilde{\theta}_t}^{\infty} b_{S,t} dF_\theta + G_t + \frac{\omega}{2}(\pi_t - 1)^2 = Y_t \tag{7}$$

$$\mathbf{1}_{\widetilde{\theta}_t<\overline{\theta}}\left[\int_{\widetilde{\theta}_t}^{\infty} \theta dF_\theta + b_{S,t} F_{\widetilde{\theta}_t} E_t \frac{u'(C_{t+1})}{\pi_{t+1}}\beta\right] + \mathbf{1}_{\widetilde{\theta}_t\geq\overline{\theta}} b_{t,S} E_t \frac{u'(C_{t+1})}{\pi_{t+1}} + b_{L,t} \sum_{j\geq 1} \beta^j \delta^{j-1} E_t \frac{u'(C_{t+j})}{\Pi_{k=1}^j \pi_{t+k}}$$
$$= \frac{b_{t-1,S}}{\pi_t} u'(C_t) + \frac{b_{t-1,L}}{\pi_t}(u'(C_t) + \delta \sum_{j\geq 1} \beta^j \delta^{j-1} E_t \frac{u'(C_{t+j})}{\Pi_{k=1}^j \pi_{t+k}}) + G_t u'(C_t) - \frac{\tau_t}{1-\tau_t} \chi Y_t^{1+\gamma} \tag{8}$$

$$\pi_t(\pi_t - 1)u'(C_t) = \frac{\eta}{\omega}\left(\frac{1+\eta}{\eta} u'(C_t) - \frac{1}{1-\tau_t} \chi Y_t^{\gamma}\right) Y_t + \beta E_t u'(C_{t+1}) \pi_{t+1}(\pi_{t+1} - 1) \tag{9}$$

and condition $\widetilde{\theta}_t C_t = b_{S,t}$.

Proposition 1 defines the set of sufficient conditions as constraints (7) to (9) and the threshold condition $\widetilde{\theta}_t v'(b_{S,t}) = u'(C_t)$. (8) is the government's budget constraint, obtained by substituting out prices and scaling by $u'(C_t)$. We use the same scaling in the Phillips curve constraint in (9).

2.2.2 Optimality

As in standard Ramsey problems with commitment, equations (8) and (9) imply that future optimal policies enter today's constraints, restricting the set of feasible policies at the current date. As a consequence, the planner's problem can be solved either in its sequential form using a Lagrangian, or recursively, following Marcet and Marimon (2019), by augmenting the state space with co-state variables that summarize past commitments.

The expressions derived for equations (8), (9), and the planner's objective, however, reveal a difficulty with the first approach in our context. When $\overline{\theta}$ is finite, the derivative of the Lagrangian with respect to $\widetilde{\theta}_t$ may not be defined.[10]

For some of our analytical examples below, we assume that F has finite support. To deal with the resulting non-differentiability of the integrals at $\overline{\theta}$, we restrict the planner's program by imposing $\widetilde{\theta}_t \leq \overline{\theta}$. For the specific examples considered, this restriction is justified, as the planner's objective is flat for $\widetilde{\theta}_t \geq \overline{\theta}$.

Setting this issue aside for the moment, we introduce the solution to the Ramsey program through the FONCs of the Lagrangian under $\overline{\theta} = \infty$. Letting $\psi_{gov,t}, \psi_{RC,t}, \psi_{PC,t}, \psi_{\widetilde{\theta},t}$ denote the multipliers attached to the government budget, the resource constraint, the Phillips curve and the constraint $C_t\widetilde{\theta}_t = b_{S,t}$ respectively, we can write the Lagrangian function as follows:

$$\begin{aligned}
\mathcal{L} = E_0 \sum_{t\geq 0} \beta^t \Bigg\{ & \widetilde{U}_t(C_t, b_{S,t}, Y_t, \widetilde{\theta}_t) + \psi_{\widetilde{\theta},t}\left(\widetilde{\theta}_t C_t - b_{S,t}\right) \\
& + \psi_{RC,t}\left(C_t + C_t \int_0^{\widetilde{\theta}_t} \theta dF_\theta + \int_{\widetilde{\theta}_t}^{\infty} b_{S,t} dF_\theta + G_t + \frac{\omega}{2}(\pi_t - 1)^2 - Y_t\right) + \\
& + \psi_{PC,t}\left(\pi_t(\pi_t - 1)u'(C_t) - \frac{\eta}{\omega}\left(\frac{1+\eta}{\eta}u'(C_t) - \frac{1}{1-\tau_t}\chi Y_t^{\gamma}\right)Y_t - \beta E_t u'(C_{t+1})\pi_{t+1}(\pi_{t+1} - 1)\right) + \\
& \psi_{gov,t}\Bigg(b_{S,t}\left[\int_{\widetilde{\theta}_t}^{\infty} \theta v'(b_{S,t}) dF_\theta + F_{\widetilde{\theta}_t} E_t \frac{u'(C_{t+1})}{\pi_{t+1}}\beta\right] + b_{L,t}\sum_{j\geq 1}\beta^j \delta^{j-1} E_t \frac{u'(C_{t+1})}{\Pi_{k=1}^{j}\pi_{t+k}} \\
& - \frac{b_{t-1,S}}{\pi_t}u'(C_t) - \frac{b_{L,t-1}}{\pi_t}\left(u'(C_t) + \delta\sum_{j\geq 1}\beta^j\delta^{j-1}E_t\frac{u'(C_{t+j})}{\Pi_{k=1}^{j}\pi_{t+k}}\right) - G_t u'(C_t) + \frac{\tau_t}{1-\tau_t}\chi Y_t^{1+\gamma}\Bigg)\Bigg\}
\end{aligned}$$

[10] This occurs whenever $f_{\overline{\theta}} \neq 0$. To see this explicitly, consider differentiating one of the terms, for example $\int_0^{\min\{\widetilde{\theta}_t, \overline{\theta}\}} \theta \, dF_\theta$. At $\widetilde{\theta}_t = \overline{\theta}$ the left derivative equals $\overline{\theta} f_{\overline{\theta}}$, whereas the right derivative equals zero.

Clearly, non-differentiability matters only if $\widetilde{\theta}_t = \overline{\theta}$ is a critical point. Elsewhere, the Lagrangian is differentiable including for $\widetilde{\theta}_t > \overline{\theta}$, in which case the constraint $C_t\widetilde{\theta}_t = b_{S,t}$ does not influence the solution.

This issue is specific to the sequential Lagrangian formulation. The alternative approach of characterizing optimal policy using the SPFE, is well defined in the presence of kinks, as optimal policies are obtained through global maximization. We derive the SPFE in Appendix D. In the main text, we adopt the Lagrangian approach, as it provides a more transparent mapping between FOCs and constraints.

The first-order conditions for the optimum are:

$$(10)\quad \widetilde{\boldsymbol{\theta}}_{\boldsymbol{t}}: \quad \overbrace{\frac{d\widetilde{U}_t}{d\widetilde{\theta}_t}+\psi_{RC,t}\left(\widetilde{\theta}_t f_{\widetilde{\theta}_t}C_t-b_{S,t}f_{\widetilde{\theta}_t}\right)}^{=0}+\psi_{\widetilde{\theta},t}C_t+\psi_{gov,t}b_{S,t}\left(\beta f_{\widetilde{\theta}_t}E_t\frac{u'(C_{t+1})}{\pi_{t+1}}-\widetilde{\theta}_t f_{\widetilde{\theta}_t}\frac{1}{b_{S,t}}\right)=0$$

$$(11)\quad \boldsymbol{Y_t}: \quad \frac{d\widetilde{U}_t}{dY_t}-\psi_{RC,t}+\psi_{PC,t}\left(\frac{1+\gamma}{1-\tau_t}\frac{\eta}{\omega}\chi Y_t^{\gamma}-\frac{1+\eta}{\omega}u'(C_t)\right)+\psi_{gov,t}\frac{\tau_t}{1-\tau_t}(1+\gamma)\chi Y_t^{\gamma}=0$$

$$(12)\quad \boldsymbol{\tau_t}: \quad \psi_{PC,t}\frac{\eta}{\omega}+\psi_{gov,t}=0$$

$$(13)\quad \boldsymbol{C_t}: \quad \frac{d\widetilde{U}_t}{dC_t}+\psi_{RC,t}\left(1+\int_0^{\widetilde{\theta}_t}\theta dF_\theta\right)+u''(C_t)\pi_t(\pi_t-1)\Delta\psi_{PC,t}-\psi_{PC,t}\frac{1+\eta}{\omega}u''(C_t)Y_t+\psi_{\widetilde{\theta},t}\widetilde{\theta}_t$$
$$-\psi_{gov,t}u''(C_t)G_t-\frac{b_{t-1,S}}{\pi_t}u''(C_t)\left(\psi_{gov,t}-F_{\widetilde{\theta}_{t-1}}\psi_{gov,t-1}\right)-u''(C_t)\sum_{j\geq 0}\delta^j\frac{b_{L,t-j-1}}{\Pi_{k=0}^{j}\pi_{t-k}}\Delta\psi_{gov,t-j}=0$$

$$\boldsymbol{\pi_t}: \ \omega\psi_{RC,t}(\pi_t-1)+u'(C_t)(2\pi_t-1)\Delta\psi_{PC,t}+u'(C_t)(1+\delta q_{L,t})\sum_{j\geq 0}\delta^j\frac{b_{L,t-j-1}}{\pi_t\Pi_{k=0}^{j}\pi_{t-k}}\Delta\psi_{gov,t-j}$$

$$(14)\quad +\frac{b_{S,t-1}}{\pi_t^2}u'(C_t)(\psi_{gov,t}-\psi_{gov,t-1}F_{\widetilde{\theta}_{t-1}})=0$$

$$(15)\quad \boldsymbol{b_{S,t}}: \quad \frac{d\widetilde{U}_t}{db_{S,t}}+\psi_{RC,t}\left(1-F_{\widetilde{\theta}_t}\right)-\psi_{\widetilde{\theta},t}+\psi_{gov,t}\beta F_{\widetilde{\theta}_t}E_t\frac{u'(C_{t+1})}{\pi_{t+1}}=\beta E_t\frac{u'(C_{t+1})}{\pi_{t+1}}\psi_{gov,t+1}$$

$$(16)\quad \boldsymbol{b_{L,t}}: \quad \psi_{gov,t}\sum_{j\geq 1}\beta^j\delta^{j-1}E_t\frac{u'(C_{t+j})}{\Pi_{k=1}^{j}\pi_{t+k}}=E_t\psi_{gov,t+1}\sum_{j\geq 1}\beta^j\delta^{j-1}\frac{u'(C_{t+j})}{\Pi_{k=1}^{j}\pi_{t+k}}$$

The first-order conditions, together with the constraints of the Ramsey program, form the system of equations that must be solved to determine the optimal policy. We next examine a subset of these conditions to highlight the key policy margins and to compare them with their counterparts in the canonical Ramsey model. We define the latter as the case in which short-term government bonds do not provide any liquidity or safety. In the expressions below, this corresponds to imposing $F_{\widetilde{\theta}_t}=1$, $f_{\widetilde{\theta}_t}=\psi_{\widetilde{\theta}_t}=\frac{d\widetilde{U}_t}{db_{S,t}}=\frac{d\widetilde{U}_t}{d\widetilde{\theta}_t}=0$.

Inflation FONC. Consider first the optimality condition for inflation (14). According to this expression, when determining the optimal inflation rate, the planner takes into account the associated social costs (measured by the resource-constraint multiplier, $\psi_{RC,t}<0$), and the effects through the Phillips curve and the government budget constraint.

To interpret the term $u'(C_t)(2\pi_t-1)\Delta\psi_{PC,t}$ (the Phillips-curve effect), suppose that the planner commits in t to a higher inflation rate at the current realized state. This commitment generates greater inflationary pressure in $t-1$, altering the feasible set of policies in that earlier period. The lagged multiplier $\psi_{PC,t-1}$ captures this effect.

Notice that this interpretation is consistent with the standard result that $\psi_{PC,t-1}$ acts as a penalty term enforcing past commitments (Marcet and Marimon, 2019). Monopolistic competition creates a wedge between the real wage and the households' marginal rate of substitution between consumption and leisure, thereby generating an inflationary bias. Away from the efficient allocation—when taxes are not sufficiently negative to fully offset the markup distortion—the planner has an incentive to

raise inflation and move the economy closer to the efficient outcome. The presence of $\psi_{PC,t-1}$ removes this incentive by penalizing deviations from previously promised policies.

An analogous interpretation applies to the terms in (14) featuring lags of the Lagrange multiplier associated with the government budget constraint. Focus on the first of these terms, which is related to long-term debt. Since this debt is modeled as a perpetuity, promises about π_t affect the prices of long bonds issued in earlier periods. For example, a promise made in period $t-j-1$ to raise inflation in t reduces the price of debt issued in that period proportionally to $\delta^j b_{L,t-j-1}$, and lowers its real value outstanding in period $t-j$ by the same factor.

The terms $\Delta\psi_{gov,t-j}$ summarize the impact of such promises on the government budget constraints in $t-j-1$ and $t-j$. Moreover, they act as penalty terms that align the ultimate inflation choice in period t with the planner's past commitments. Without them, the planner would have an incentive to inflate away outstanding long-term debt, generating an inflationary debt bias. The lagged multipliers internalize this intertemporal tradeoff and eliminate that incentive.[11]

Finally, the term in the bottom line of (14) captures analogous channels operating through short-term bond prices. A promise of higher inflation in t, made in $t-1$, affects the issuance price of short debt through the component $\beta F_{\widetilde{\theta}_{t-1}} E_{t-1} \frac{u'(C_t)}{u'(C_{t-1})\pi_t} b_{S,t-1}$. In addition, higher π_t reduces the real repayment value of government debt outstanding from $t-1$.

Contrast (14) with the analogous expression that would prevail in the canonical Ramsey model, when short-term debt does not provide liquidity/safety to the private sector. Simply set $F_{\widetilde{\theta}_t} = 1$ and the expression for optimal inflation remains otherwise unchanged. The new term $\frac{b_{S,t-1}}{\pi_t} u'(C_t)(\psi_{gov,t} - \psi_{gov,t-1})$ has an interpretation similar to before: it captures how inflation affects consecutive government budget constraints and enforces consistency with past commitments. The key difference, however, is that a promise of higher inflation in the current state now affects the issuance price of debt in $t-1$ by $\beta E_{t-1} \frac{u'(C_t)}{u'(C_{t-1})\pi_t}$, instead of $F_{\widetilde{\theta}_{t-1}} \beta E_{t-1} \frac{u'(C_t)}{u'(C_{t-1})\pi_t}$ as in the baseline model.

At the center of this discrepancy is the notion that when debt provides liquidity/safety to agents, they do not require full compensation for higher expected inflation. Another way of reading the 'penalty' term $F_{\widetilde{\theta}_{t-1}} \psi_{gov,t-1}$ is the following: When $F_{\widetilde{\theta}_{t-1}}$ is less than one, the lagged multiplier $\psi_{gov,t-1}$ still constrains current inflation, but its effect is attenuated. Economically, this resembles a partial-commitment environment: the planner is constrained by past promises, yet the commitment is weaker than in the canonical Ramsey case. In the extreme, if $F_{\widetilde{\theta}_{t-1}} \approx 0$, the lagged multiplier no longer affects current inflation, and the planner behaves as if there were no commitment, even though the underlying mechanism here differs from discretion.

Optimal bond supplies. Equations (15) and (16) are the first-order conditions for short- and long-term bonds, respectively.

According to (16), the multiplier on the government budget constraint, $\psi_{gov,t}$, follows a risk-

[11] Note that (12), the first-order condition for τ_t, implies that the planner sets the tax rate so that the shadow cost of debt is proportional to the shadow cost of inflation. Thus, if the economy is operating at the efficient allocation, then $\psi_{PC,t} = \psi_{gov,t} = 0$. Inflation is zero even when the planner is not constrained by past commitments (i.e., when $\psi_{gov,t-1} = \psi_{PC,t-1} = 0$), reflecting that there is no incentive to inflate away debt or push output above potential.

adjusted random walk: $\psi_{gov,t} = \frac{E_t \psi_{gov,t+1} \varpi_{t+1}}{E_t \varpi_{t+1}}$ with risk measure $\varpi_{t+1} \equiv \sum_{j \geq 1} \beta^j \delta^{j-1} E_{t+1} \frac{u'(C_{t+j})}{\Pi_{k=1}^{j} \pi_{t+k}}$.

This property is standard in models of optimal debt policy under incomplete financial markets (see Aiyagari et al., 2002). The multiplier measures the marginal cost of raising public funds by issuing non-contingent long debt to smooth taxes intertemporally. The martingale property implies that this cost is smoothed over time. Following a positive spending shock that must be financed with higher taxes, the optimal policy induces a permanent increase in $\psi_{gov,t}$ reflecting the planner's objective to spread the tax distortions as evenly as possible across time.

Equation (15), however, shows that an analogous property does not generally characterize the optimal policy for short-term debt. First, the bond supply directly affects the welfare objective function, the resource constraint, and the threshold $\widetilde{\theta}_t$, so that the first three terms in (15) are generically nonzero. Second, the last two terms in (15) are also inconsistent with a random-walk when $F_{\widetilde{\theta}_t} < 1$.

To interpret this condition, we follow the discussion in Angeletos et al. (2022). For convenience, (15) can be rewritten as:

$$\underbrace{\frac{d\widetilde{U}_t}{db_{S,t}} + \psi_{RC,t}\left(1 - F_{\widetilde{\theta}_t}\right) - \psi_{\widetilde{\theta},t}}_{\text{Social Benefit}} + \underbrace{\psi_{gov,t}\beta(F_{\widetilde{\theta}_t} - 1)E_t\frac{u'(C_{t+1})}{\pi_{t+1}}}_{\text{Liquidity/Safety Rent}} = \beta \underbrace{E_t\frac{u'(C_{t+1})}{\pi_{t+1}}\Delta\psi_{gov,t+1}}_{\text{Tax Smoothing}} \tag{17}$$

Equation (17) can be interpreted as balancing the social value of providing the liquid/safe asset—namely, the gain from easing the financial friction net of the rent-suppression cost—against the welfare cost of departing from tax smoothing (as encoded in the random-walk property of the Lagrange multiplier on the RHS). The first terms of the LHS capture the social benefit of liquidity/safety provision (in terms of direct utility impact, the aggregate demand effect and the impact on the threshold constraint). The second term is the marginal effect of increasing $b_{S,t}$ on the rent extracted. Overall, (17) implies that the planner optimally departs from tax smoothing to account for these effects.

Angeletos et al. (2022) explain that the process defined in (17) (the analogue in their model) induces transitional dynamics in the Lagrange multiplier towards a steady state in which the financial friction binds. Differently from their paper, however, in our model not all of the government debt is liquid/safe. Equation (17) codetermines the path of the multiplier together with (16) and, as we will later see, our model does not admit any transitional dynamics in $\psi_{gov,t}$. Intuitively, the planner can instantaneously finance the desired level of short-term debt, by borrowing (or lending) long term. This mechanism will be at the heart of the optimal debt management exercises that we will conduct.

Finally, it is straightforward to show that in the case where $F_{\widetilde{\theta}_t} = 1$, all LHS terms in (17) vanish, which restores the random walk property of the multiplier in the short bond Euler equation.

Consumption FONC. Equation (13) shows that the optimal path of C_t is governed by forces analogous to those that determine optimal inflation. Combining (10) with (13) and substituting

the multiplier $\psi_{\widetilde{\theta},t}$ —internalizing the effect of consumption through $\widetilde{\theta}_t$—we obtain the following expression:

$$
\begin{aligned}
0 = \frac{d\widetilde{U}_t}{dC_t} + \psi_{RC,t}\left(1 + \int_0^{\widetilde{\theta}_t} \theta dF_\theta\right) - \psi_{PC,t}\frac{1+\eta}{\omega}u''(C_t)Y_t - \psi_{gov,t}u''(C_t)G_t \\
\overbrace{-\widetilde{\theta}_t^2\psi_{gov,t}f_{\widetilde{\theta}_t}u'(C_t)\left(\beta E_t\frac{u'(C_{t+1})}{u'(C_t)\pi_{t+1}} - 1\right)}^{\text{Liquidity Cutoff}} + \overbrace{u''(C_t)\pi_t(\pi_t - 1)\Delta\psi_{PC,t}}^{\text{Phillips curve discount}} \\
\underbrace{-\frac{b_{t-1,S}}{\pi_t}u''(C_t)\left(\psi_{gov,t} - F_{\widetilde{\theta}_{t-1}}\psi_{gov,t-1}\right)}_{\text{Short-term real rate}} \underbrace{- u''(C_t)\sum_{j\geq 0}\delta^j\frac{b_{L,t-j-1}}{\Pi_{k=0}^j\pi_{t-k}}\Delta\psi_{gov,t-j}}_{\text{Long-term real rate}}
\end{aligned}
\tag{18}
$$

All terms featuring changes in $\psi_{PC,t}$ and $\psi_{gov,t}$ can be interpreted analogously to the corresponding terms in the inflation FONC.[12] Intuitively, consumption commitments affect feasibility through their impact on real interest and discount rates, which in turn alter the government's borrowing costs and firms' profit valuations (and thus pricing decisions). Hence, for example, the term $\frac{b_{t-1,S}}{\pi_t}u''(C_t)\left(\psi_{gov,t} - F_{\widetilde{\theta}_{t-1}}\psi_{gov,t-1}\right)$ encapsulates the planner's promise to reduce consumption in period t (relative to C_{t-1}), which raises the short bond price at issuance in $t-1$. As in the case of higher inflation, agents are willing to hold short-term debt despite a lower real return. Analogous arguments can be utilized to explain the remaining terms featuring first differences of multipliers.

Finally, the term labeled 'Liquidity cutoff' in (18) captures how current consumption affects the equilibrium ratio $\widetilde{\theta}_t$. A marginal decrease in this ratio (equivalently, an increase in C_t) reallocates a mass $f_{\widetilde{\theta}_t}$ across the cutoff, from the constrained region to the unconstrained region. For these agents, the marginal benefit from holding short-term debt becomes $\beta E_t\frac{u'(C_{t+1})}{u'(C_t)\pi_{t+1}}$ instead of its liquidity value which (in marginal utility terms) is normalized to 1. The planner internalizes this effect through the government budget constraint.

In a subsequent section, we will return to the formula (18) and use it to derive implications for optimal taxation. Then we will discuss more extensively how the above margins impact optimal fiscal policy.

2.3 Deterministic Models/ Long Run Steady State

We now turn to the characterization of the steady-state solution to the optimality conditions. We focus on a subset of these conditions to highlight key properties of optimal inflation and the maturity

[12] Since the terms in the top line of (18) correspond to standard marginal effects of consumption on welfare, aggregate demand, marginal costs, and the government's surplus, we leave them unlabeled in the equation.

composition of debt. Evaluating (14), (15) and (16) in steady state, we obtain:[13]

$$\frac{d\widetilde{U}}{d\bar{b}_S} + \psi_{RC}\left(1 - F_{\widetilde{\theta}}\right) - \psi_{\widetilde{\theta}} + \beta\frac{u'(\overline{C})}{\overline{\pi}}\psi_{gov}\left(F_{\widetilde{\theta}} - 1\right) = 0 \tag{19}$$

$$\psi_{gov}\sum_{j\geq 1}\beta^j\delta^{j-1}\frac{u'(\overline{C})}{\overline{\pi}^j} = \psi_{gov}\sum_{j\geq 1}\beta^j\delta^{j-1}\frac{u'(\overline{C})}{\overline{\pi}^j} \tag{20}$$

$$-\omega\psi_{RC}(\overline{\pi}-1)\overline{\pi}^2 = u'(\overline{C})\bar{b}_S\psi_{gov}(1 - F_{\widetilde{\overline{\theta}}}) \tag{21}$$

Consider first the optimal inflation, as determined by (21). The following property follows directly from this condition.

Proposition 2 (Optimal inflation rate) *In a steady state satisfying $F_{\widetilde{\theta}} < 1$ and $\psi_{gov} > 0$, the optimal inflation rate is strictly positive, that is, $\overline{\pi} > 1$.*

This follows immediately from (21) noting that $\psi_{RC} < 0$ and all RHS terms are positive. Hence, $(\overline{\pi}-1)\overline{\pi}^2 > 0$ requires $\overline{\pi} > 1$.

The previous discussion hinted at this property. In an equilibrium with binding frictions- when agents value government debt for its liquidity/safety-they are willing to hold it even if inflation erodes its real return. The planner exploits this margin by setting a positive inflation target. Intuitively, cheaper debt financing reduces the required primary surplus, thereby lowering the distortionary tax rate in steady state. The planner trades off inflation distortions against lower tax distortions.

Equation (21) determines a positive inflation target only when the ratio of the face value of short debt over consumption is low enough to have $F_{\widetilde{\overline{\theta}}} < 1$. Obviously, if the supply of short bonds is such that $F_{\widetilde{\theta}} = 1$, the unique solution to (21) is $\overline{\pi} = 1$.

Thus, determining whether inflation is zero or positive in the steady state equilibrium requires solving for the optimal portfolio. However, equations (19), (20) are not informative in this regard. To see this, consider first the optimality condition for long-term bonds, (20). This equation holds trivially in steady state for any $\bar{b}_L$. Therefore, it does not pin down a unique optimal level of long-term debt, a standard property of the canonical Ramsey model.[14]

Next, consider (19). A unique solution is not guaranteed here either. In particular, there may be multiple solutions satisfying $F_{\widetilde{\theta}} < 1$. Moreover, multiplicity is certain if the support of F is bounded above such that $F_{\widetilde{\theta}} = 1$ for all $\widetilde{\theta} \geq \overline{\theta}$ and so (19) again holds trivially in steady state.

For our subsequent analysis, we focus on cases in which a unique solution satisfying $F_{\widetilde{\theta}} < 1$ coexists with multiple critical points at which $F_{\widetilde{\theta}} = 1$. In such situations—relevant for the calibrations

[13] We use bars to denote the steady-state values of the variables, except for the multipliers and $\widetilde{\theta}$.

[14] The indeterminacy in the steady state of the canonical model is well known. In this case, the multiplicity of solutions is not problematic as portfolios have no bearing on the optimal allocation and therefore on welfare. Analogously, in our model for any vector $(\bar{b}_S, \bar{b}_L)$ such that $F_{\widetilde{\theta}} = 1$ the optimal policy solution for the remaining variables is identical and therefore the welfare function is constant over $(\bar{b}_S, \bar{b}_L)$. However, when such solutions coexist with a steady state with binding financial frictions, then determining the optimal policy is not trivial as we discuss next.

below—the system of first-order conditions is not sufficient to pin down the equilibrium. Instead, the global optimum must be identified by evaluating the welfare function at each stationary point. Long-term debt can be chosen so as to select the optimum and (20) is guaranteed to hold.[15]

In such situations—relevant for the calibrations below—the optimum cannot be determined solely from the system of first-order conditions. Instead, identifying the global optimum requires evaluating the welfare function at each of the stationary points. Long-term debt can be chosen so as to select the optimum and (20) is guaranteed to hold.[16]

2.3.1 Global optimum under quasilinear preferences

We now present an analytical example of the globally optimal Ramsey policy under quasilinear preferences and fully rigid prices. For tractability, we assume that f is uniform in $[0, \overline{\theta}]$. Rather than restricting attention to the long run steady state, we solve the full Ramsey program (initiating at date 0) under no uncertainty. As we demonstrate, under the assumptions of this paragraph, the Ramsey program reduces to a static optimization problem for which convergence to the steady state is instantaneous and an analytical solution of the optimal policy is feasible. Demonstrating this result is useful, since it highlights key properties of the model that will play an important role in the analysis that follows.

Our main finding, stated in Proposition 3 below, is that in this environment it is optimal to set $\widetilde{\theta}_t < \overline{\theta}$ whenever the debt constraint is relevant for the optimal policy. By contrast, when $\psi_{gov} = 0$ (at the first best) then $\widetilde{\theta}_t \geq \overline{\theta}$.

The full Ramsey program under quasilinear utility/rigid prices. Under the assumptions of this section, we can express the Ramsey program as (see appendix):

$$\max_{\{Y_t, C_t, \widetilde{\theta}_t, b_{S,t}, b_{L,t}\}} \sum_{t\geq 0} \beta^t \left[C_t + \mathcal{I}_{\widetilde{\theta}_t < \overline{\theta}} \left(\int_0^{\widetilde{\theta}_t} \theta \log(\theta) \frac{1}{\overline{\theta}} d\theta + \int_{\widetilde{\theta}_t}^{\overline{\theta}} \theta \log(\widetilde{\theta}_t) \frac{1}{\overline{\theta}} d\theta \right) + \mathcal{I}_{\widetilde{\theta}_t \geq \overline{\theta}} \int_0^{\overline{\theta}} \log(\theta) \frac{1}{\overline{\theta}} d\theta - \chi \frac{Y_t^{1+\gamma}}{1+\gamma} \right]$$

$$\text{s.t} \quad \sum_{j\geq 0} \beta^j \left[\left(\int_{\widetilde{\theta}_{t+j}}^{\overline{\theta}} \frac{\theta}{\overline{\theta}} d\theta + \beta \left(\frac{\widetilde{\theta}_{t+j}}{\overline{\theta}} - 1 \right) b_{S,t+j} \right) \mathcal{I}_{\widetilde{\theta}_{t+j} < \overline{\theta}} + Y_{t+j}(1 - \chi Y_{t+j}^{\gamma}) - \overline{G} \right]$$

$$= b_{t-1,S} + b_{t-1,L} \sum_{j\geq 0} \delta^j \beta^j \qquad t = 0, 1, 2, ... \tag{22}$$

$$b_{S,t} = \widetilde{\theta}_t \tag{23}$$

$$C_t + \overline{G} + \mathcal{I}_{\widetilde{\theta}_t < \overline{\theta}} \left(\int_0^{\widetilde{\theta}_t} \theta \frac{1}{\overline{\theta}} d\theta + \widetilde{\theta}_t \int_{\widetilde{\theta}_t}^{\overline{\theta}} \frac{1}{\overline{\theta}} d\theta \right) + \mathcal{I}_{\widetilde{\theta}_t \geq \overline{\theta}} \left(\int_0^{\overline{\theta}} \theta \frac{1}{\overline{\theta}} d\theta \right) = Y_t \tag{24}$$

given $b_{-1,S}, b_{-1,L}$. Equation (23) is the analogue of the threshold constraint in this model. (24) defines the sequence of resource constraints. Finally, using (22), we have replaced the sequence

[15] This is obviously a steady-state comparison and therefore a static implementation argument. In the dynamic setting, substitutions between short and long-term debt may not be trivial.

[16] This is obviously a steady-state comparison and therefore a static implementation argument. In the dynamic setting, substitutions between short and long-term debt may not be trivial.

of per-period budget constraints with the intertemporal constraints for $t = 0, 1, 2, ...$ This is an equivalent formulation.

The next lemma shows that this Ramsey program can be simplified into a static optimization problem.

Lemma 1 *Consider assumptions $u(C_t) = C_t$, $\frac{1+\eta}{\eta} = 1$ and $G_t = \overline{G}$ for all t. Moreover, suppose that f is uniform in $[0, \overline{\theta}]$. The optimal Ramsey policy can be found as a solution to the following static optimization problem.*

$$\max_{\overline{Y}, \widetilde{\theta}} \overline{Y} - \frac{1}{\overline{\theta}} \left(\frac{1}{2} \widetilde{\theta}^2 + \widetilde{\theta}(\overline{\theta} - \widetilde{\theta}) \right) + \frac{\overline{\theta}}{2} \log(\widetilde{\theta}) - \frac{\widetilde{\theta}^2}{4\overline{\theta}} - \chi \frac{\overline{Y}^{1+\gamma}}{1+\gamma}$$

$$\textit{subject to} \quad \widetilde{\theta} \leq \overline{\theta} \qquad \textit{and} \qquad \Upsilon = \underbrace{\beta(\frac{\widetilde{\theta}}{\overline{\theta}} - 1)\widetilde{\theta} + \frac{1}{2\overline{\theta}} \left(\overline{\theta}^2 - \widetilde{\theta}^2 \right)}_{\textit{Rents} \geq 0 \quad (=0 \textit{ if } \widetilde{\theta} = \overline{\theta})} + \overline{Y}(1 - \chi \overline{Y}^{\gamma})$$

$\Upsilon \equiv (1-\beta)B_{-1} + \overline{G}$ and B_{-1} denotes total initial debt (both long and short).

Although the appendix proves this result, it is useful to briefly discuss the intuition behind it. In particular, three points are worth making: First, under no uncertainty, the constraints (22) for $t > 0$ are redundant and therefore can be omitted from the maximization. Intuitively, for a given optimal allocation $\widetilde{\theta}_t, C_t, Y_t, b_{S,t}$ the planner can use long-term debt to ensure satisfaction of the date $t > 0$ constraints. Second, for any $\widetilde{\theta}_t \geq \overline{\theta}$ the optimal quantity $\widetilde{\theta}_t$ is not defined in the resulting program (the welfare objective is flat). Hence, we can let the planner choose $\widetilde{\theta}_t$ subject to the inequality constraint $\widetilde{\theta}_t \leq \overline{\theta}$, and the corner solution fully characterizes optima in which short-term debt does not provide liquidity/safety services. The integrals appearing in the Ramsey program are well defined and so are their (left) derivatives with respect to $\widetilde{\theta}_t$.

Third, the assumptions of quasilinear utility and fully rigid prices are crucial for reducing the full dynamic Ramsey problem to the static formulation presented in Lemma 1. Under these assumptions, the optimal allocation is time invariant. The planner has no incentive to distort prices intertemporally to relax the constraint (22). Finally, note that the assumptions that f is uniform and the absence of monopoly distortions are not consequential for our results in this paragraph.

Our next result characterizes the solution to this Ramsey program.

Proposition 3 (Optimal cut-off) *The solution to the program defined in Lemma 1 satisfies $\widetilde{\theta}^* < \overline{\theta}$ when $\psi_{gov} > 0$ (interior optimum) and $\widetilde{\theta}^* = \overline{\theta}$ when $\psi_{gov} = 0$ (corner solution).*

Proposition 3 shows that whenever the government debt constraint is binding, the planner limits the supply of the short-term asset, thereby extracting liquidity/safety rents. Positive rents slacken the constraint enabling to finance debt at lower taxes. At the first best allocation, $\psi_{gov} = 0$ and this fiscal motive for restricting the short-term asset disappears. Setting $\widetilde{\theta} = \overline{\theta}$ becomes optimal.[17]

[17] It is simple to show that for $\psi_{gov} < 0$ (or initial debt is negative enough so that increasing it raises welfare by moving the economy closer to the first best), the optimal policy is also the corner solution $\widetilde{\theta} = \overline{\theta}$.

Comparative Statics. The static program in Proposition 3 yields a unique interior solution (see appendix C). The following result reveals how the solution changes with the initial debt and spending levels, Υ.

Corollary 3.1 *Let $\widetilde{\theta}^*(\Upsilon)$ denote the interior optimum. It holds that $\frac{d\widetilde{\theta}^*(\Upsilon)}{d\Upsilon} < 0$.*

Thus, assuming higher spending or initial debt, leads the planner to target a lower equilibrium $\widetilde{\theta}$.

Comparison with Angeletos et al. (2022). Under quasilinear utility and rigid prices, the structure of our Ramsey program is similar to the policy problem analyzed in Angeletos et al. (2022). In that paper, higher initial debt may lead to an equilibrium in which the economy is satiated with liquidity (see their Theorem 1). We found the opposite; higher initial debt always translates to a lower target $\widetilde{\theta}^*$. The canonical outcome obtains when $\psi_{gov} = 0$.

In contrast to Angeletos et al. (2022), in our setting, not all government debt provides liquidity/safety. Faced with a high initial burden, the planner immediately reduces the supply of short bonds and increases the supply of long-term debt. By contrast, if only liquid/safe short bonds can be issued (or if we constrain the quantity of long bonds not to exceed a certain value, see below), one cannot dispense with the constraints (22) for $t \geq 1$, and immediate convergence to the steady state need not obtain. In that case, it becomes necessary to explicitly characterize the evolution of the multiplier $\psi_{gov,t}$, which is given by:

$$\psi_{gov,t+1} = \begin{cases} \psi_{gov,t}\dfrac{\widetilde{\theta}_t}{\overline{\theta}}\left(2 - \dfrac{1}{\beta}\right) + \dfrac{(\overline{\theta} - \widetilde{\theta}_t)^2}{2\beta\widetilde{\theta}_t\overline{\theta}}, & \text{if } \widetilde{\theta}_t \leq \overline{\theta}, \\ \psi_{gov,t}, & \text{if } \widetilde{\theta}_t \geq \overline{\theta}. \end{cases} \tag{25}$$

Following the argument of Angeletos et al. (2022), we can show that if the economy starts with a high enough initial debt, then the cost of distorting taxes intertemporally, to reach a steady state in which $\widetilde{\theta} < \overline{\theta}$, is too high. The planner sets $\widetilde{\theta}_t > \overline{\theta}$ for all t and the standard Ramsey outcome applies. By contrast, for sufficiently low initial debt, the planner instead targets the steady state characterized in Proposition 3. The economy exhibits transitional dynamics, as captured by (25).

These considerations are absent here, as issuing long-term government debt ensures that $\psi_{gov,t} = \psi_{gov,t+1}$ for all t.

2.3.2 Curved Utility/Sticky Prices/Multiplicity.

The previous subsection considered the solution to the Ramsey program when preferences are quasilinear and prices are rigid. Away from this constrained scenario, when utility is curved and prices are sticky, obtaining a closed form characterization of the Ramsey solution for a general initial maturity structure is not straightforward. Under curved utility, the market value of initial debt depends on the path of consumption and, in general, the planner will distort taxes intertemporally to manipulate

bond prices (e.g. Faraglia et al., 2016). Analogously, under sticky prices the planner will distort inflation so as to reduce initial debt. In both of these cases, the economy will in principle feature transitional dynamics in the endogenous policy variables, even though the instantaneous convergence of ψ_{gov} continues to hold.

Given that simple analytical solutions are not available, it is also not simple to determine whether along the transition the planner will constrain the quantity of short bonds or provide abundant liquidity to the economy. Such a characterization would require verifying in every period the global optimum evaluating (backwards) the welfare objective function. In Section 4 we will develop a numerical algorithm that is well suited to characterize such solutions.

Figure 1: **Welfare as a function of the short bond supply**.

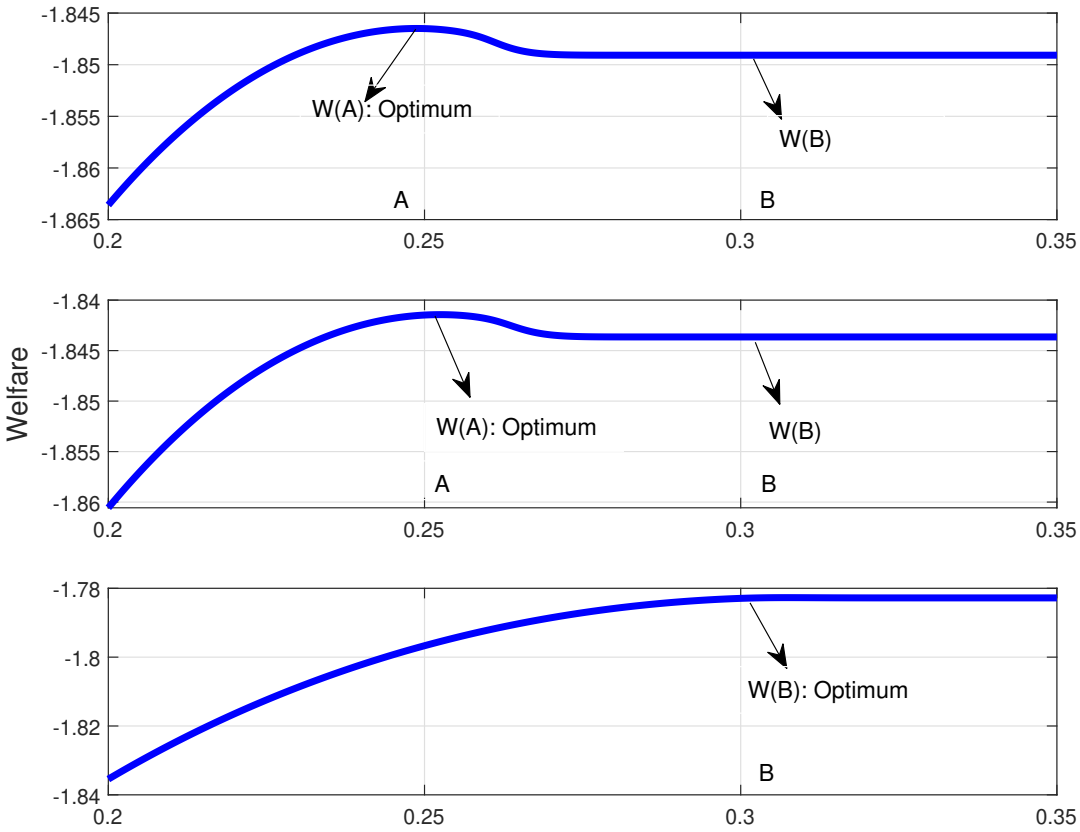

Notes: We plot the welfare objective defined in Proposition 1 as a function of the level of short-term debt. See text for details on how the three graphs have been constructed. $W(A)$ and $W(B)$ denote welfare evaluated at A and B respectively.

To conserve space—and since numerical simulations of the full model with aggregate shocks will illustrate transitional dynamics—we do not apply our algorithm here to solve for the date-0 allocation. Instead, we focus on the properties of the equilibrium from a *timeless perspective*, that is, assuming that the policymaker's commitments began in the distant past so that initial conditions are irrelevant. In the absence of shocks, the resulting long run equilibrium is stationary and coincides

with the steady state of the Ramsey system. We analyze the conditions under which the global optimum corresponds to an allocation with binding financial frictions.

Figure 1 plots the welfare objective as a function of short-term debt issuance under three parameterizations of the model. Each parameterization is associated with a different optimal steady-state level of long-term debt. The top panel corresponds to the calibration in which the optimal supply of long-term bonds is positive (our baseline, see Section 4.3). In the middle panel, the global optimum satisfies $\bar{b}_L = 0$. Finally, the bottom panel considers a case with negative long-term debt, $\bar{b}_L = -0.5$.[18] In each case, starting from the global optimum, we vary short-term debt issuance and adjust long-term debt endogenously so as to keep total debt fixed at its steady-state level, thereby tracing out the welfare function.

Using the graphs we can easily determine at which points the first-order conditions of the Ramsey program are satisfied. Focus on the top panel. The welfare function is first concave and then essentially flat at high short-term debt levels (past point B where $F_{\widetilde{\theta}} \approx 1$,).[19] The global optimum is at point A. However, the solution to the system of first-order conditions will identify several stationary points including A but also all points $\bar{b}_S \geq$B.

The middle panel exhibits a similar pattern. By contrast, allowing for negative long-term debt—as in the bottom panel—substantially alters the shape of the welfare function. In this case, the global optimum is attained at point B or any point to its right.

Table 1 (top panel) reports the optimal values of short-bond supply $\bar{b}_S$, the annualized inflation rate $\overline{\pi}$, and the steady-state tax rate $\overline{\tau}$ for each of the three optima shown in Figure 1. Consider first the case of positive long-term debt (column 1). We have $\overline{\pi} = 2.16\%$ and $\overline{\tau} = 0.211$. For comparison, the bottom panel reports the corresponding steady-state allocation in the canonical Ramsey model, evaluated at the same short- and long-term debt quantities as in the global optimum. Optimal inflation is zero and the steady-state tax rate is higher ($\overline{\tau} = 0.243$). These numbers are suggestive of the government's gains in terms of reducing tax distortions through rents, and of the steady-state tradeoff between inflation and taxes.

Columns 2 and 3 report the cases with zero and negative long-term debt, respectively. Two patterns emerge. First, lower levels of long-term debt are associated with higher equilibrium supplies of short-term debt (consistent with Corollary 3.1). Second, when long-term debt is negative, the short-bond supply reaches $\bar{b}_S = 0.308$, and the model outcomes closely resemble those of the canonical Ramsey allocation. Since $\overline{\tau} = 0.196 > \frac{1}{1+\eta}$ the economy is, however, not at the first best allocation.

[18] The calibration of the model parameters follows the baseline specification reported in Section 4.3.

[19] In our baseline model we assume that that $f_{\widetilde{\theta}}$ is log-normal, hence, technically, the support is unbounded. However, given the discipline for the moments of the distribution imposed by the data (see the calibration section), $f_{\widetilde{\theta}}$ is essentially zero after B so that a numerical solver will find a local optimum at any of these points. This is not the case when b_S is less than the level corresponding to B, even though the welfare objective is quite flat.

We discuss this issue further below and explain in detail how we set the threshold for variable $\widetilde{\theta}$, above which we consider an allocation to be effectively a canonical Ramsey outcome.

Table 1: **Optimal policies in steady state**

	Baseline Calibration	**Zero Debt**	**Negative Debt**
$\bar{b}_L$	0.096	0	-0.5
Financial Frictions			
$4(\overline{\pi}-1)$	2.16	2.08	0.00
$\overline{\tau}$	0.211	0.208	0.196
$\bar{b}_S$	0.249	0.252	0.308
Canonical Ramsey			
$4(\overline{\pi}-1)$	0.00	0.00	0.00
$\overline{\tau}$	0.243	0.238	0.196
$\bar{b}_S$	0.249	0.252	0.308

Notes: The table reports optimal steady-state inflation, tax rates, and short-bond supply under alternative long-term debt positions. The bottom panel reports these objects in the canonical Ramsey model.

2.4 Discussion and preview of global solution results

We derive two crucial findings from the analysis in this section. First, the level of total debt determines whether financial frictions bind at the optimum. When debt is sufficiently negative, the optimal policy supplies abundant liquidity. Second, solving the system of first-order necessary conditions is, in general, not sufficient to determine the global optimum.

These findings are particularly relevant for the algorithm we construct to solve the model with aggregate shocks, which complements the system of first-order conditions with explicit evaluations of the welfare objective function. Anticipating the results in Section 4, our simulations show that the steady-state insights extend to the Ramsey problem with aggregate shocks and transitional dynamics. In particular, for debt levels that are not too negative, the optimal policy consistently targets a low supply of short-term debt to extract rents.

In the next section, we establish additional analytical results regarding the properties of optimal taxes and inflation. Some of our derivations rely on a local approximation around the steady state with binding financial frictions. In conducting this local analysis, we take as given the global solution result that the optimal policy does not alternate between regimes of binding frictions and liquidity/safety abundance. In other words, we impose that the optimal policy varies smoothly and does not change discontinuously in response to small perturbations around the steady state.

3 Optimal Taxation and Inflation

We now investigate how the planner uses tax and inflation policies in the face of shocks to government spending when financial frictions are binding, and derive the following results. First, following a spending shock, the planner permanently raises inflation, so that the optimal policy is character-

ized by a time-varying inflation target. Second, by analytically characterizing the tax schedule, we illustrate that financial frictions introduce novel channels of intertemporal tax distortions.

3.1 Inflation targets

Our first result concerns the response of inflation to a government spending shock. To derive analytical formulae, we consider the model with quasilinear preferences and uniform f, as in the previous section, but do not impose rigid prices. The optimal inflation path can be characterized by combining the first-order condition for inflation with the short bond Euler equation. As shown in the Appendix C, a first-order expansion around the steady state with binding financial frictions yields the following expressions:

$$\hat{\pi}_t = \left[\frac{(2\overline{\pi}-1)}{(3\overline{\pi}-2)}\frac{\psi_{gov}}{\eta} + \frac{(\overline{\pi}-1)}{(3\overline{\pi}-2)}\frac{1}{(1-\frac{\widetilde{\theta}}{\overline{\theta}})}\right]\Delta\hat{\psi}_{gov,t} + \frac{(\overline{\pi}-1)}{(3\overline{\pi}-2)}\left[1-\frac{\frac{\widetilde{\theta}}{\overline{\theta}}}{(1-\frac{\widetilde{\theta}}{\overline{\theta}})}\right]\hat{b}_{S,t-1} + \frac{(\overline{\pi}-1)}{(3\overline{\pi}-2)}\hat{\psi}_{gov,t-1} + \frac{\overline{b}_L}{\omega(\overline{\pi}-\beta\delta)(3\overline{\pi}-2)}\sum_{j\geq 0}(\frac{\delta}{\overline{\pi}})^j\Delta\hat{\psi}_{gov,t-j} \tag{26}$$

$$-\frac{1}{\widetilde{\theta}\psi_{gov}}\hat{\widetilde{\theta}}_t\int_{\widetilde{\theta}}^{\overline{\theta}}\frac{\theta}{\overline{\theta}}d\theta - \frac{\beta}{\overline{\pi}}E_t\hat{\psi}_{gov,t+1} + \frac{\widetilde{\theta}}{\overline{\theta}}\left(2\frac{\beta}{\overline{\pi}}-1\right)\left(\hat{\psi}_{gov,t}+\hat{\widetilde{\theta}}_t\right) - \frac{\beta}{\overline{\pi}}\left[2\frac{\widetilde{\theta}}{\overline{\theta}}-1\right]E_t\hat{\pi}_{t+1} = 0 \tag{27}$$

where hats denote log deviations from steady state.

Combining (26) and (27) and using the property $E_t\hat{\psi}_{gov,t+1} = \hat{\psi}_{gov,t}$ (the long bond optimality condition in this linearized model) we find the following solution for optimal inflation:

$$\hat{\pi}_t = \overbrace{\mu_1\Delta\hat{\psi}_{gov,t}}^{\textbf{Short debt + Surplus Channels}} + \underbrace{\mu_2\hat{\psi}_{gov,t-1}}_{\textbf{Time-Varying Inflation Target}} + \underbrace{\mu_3\sum_{j\geq 0}(\frac{\delta}{\overline{\pi}})^j\Delta\hat{\psi}_{gov,t-j}}_{\textbf{Real Value of Long Debt Channel}} \tag{28}$$

where coefficients μ are characterized in the appendix C. From (28), we can conclude the following:

Proposition 4 (Time-varying inflation target) *The optimal policy in the model with financial frictions is characterized by a time-varying inflation target. In response to an unanticipated increase in government spending, the planner permanently increases inflation.*

Our discussion in the remainder of this subsection is devoted to interpreting this result. First, to understand why (28) implies that inflation increases permanently following a spending shock, note that since $\hat{\psi}_{gov,t}$ follows a random walk, the model solution implies $\Delta\hat{\psi}_{gov,t} = \nu\epsilon_{G,t}$ where $\nu > 0$ is a constant (uniquely) pinned down by the policy equilibrium conditions and $\epsilon_{G,t}$ is the spending shock.[20]

[20] Although the random walk property may in principle be consistent with multiple equilibria, the intertemporal constraint together with (27) select a unique value of ν. See the Appendix for an analytical characterization of ν.

It follows that $\hat{\psi}_{gov,t-1}$ can be written as the infinite sum of past spending shocks. Consequently, a shock occurring in period t raises inflation above its steady-state value in periods $t+1, t+2, \ldots$. The optimal inflation target shifts upward permanently.

Fiscal Channels of Inflation. To further examine the optimal policy and the remaining terms in (28), it is useful to consider the government's intertemporal budget constraint in its nonlinear form:

$$E_t \sum_{j\geq 0} \beta^j \Bigg[\underbrace{\left(\int_{\widetilde{\theta}_{t+j}}^{\overline{\theta}} \frac{\theta}{\overline{\theta}} d\theta + \beta \left(\frac{\widetilde{\theta}_{t+j}}{\overline{\theta}} - 1 \right) \frac{b_{S,t+j}}{\pi_{t+j+1}} \right)}_{\textbf{Rent}_{t+j}} + \text{Surplus}_{t+j} \Bigg] = \frac{b_{t-1,S}}{\pi_t} + b_{t-1,L} E_t \sum_{j\geq 0} \frac{\delta^j \beta^j}{\Pi_{k=0}^{j} \pi_{t+k}} \tag{29}$$

To build intuition, we start from a familiar environment, the canonical Ramsey model. Using (26) and imposing that $\overline{\pi} = 1$ and $\frac{\widetilde{\theta}}{\overline{\theta}} = 1$ we find that inflation evolves according to:

$$\hat{\pi}_t = \frac{\psi_{gov}}{\eta} \Delta \hat{\psi}_{gov,t} + \frac{\overline{b}_L}{\omega(1-\beta\delta)} \sum_{j\geq 0} \delta^j \Delta \hat{\psi}_{gov,t-j} \tag{30}$$

Thus, deviations of inflation from target are only temporary.

Each of the two terms on the right-hand side of (30) corresponds to one of the margins through which inflation affects — and contributes to the satisfaction of — the constraint (29).[21] Specifically, the second term in (30) maps into the real market value of long-term debt. Through this margin, a positive spending shock induces an initial increase in inflation, followed by a monotonic decline toward zero at rate δ. This path is optimal because it mirrors the rate at which the real value of coupon payments decays over time.[22]

The first term in (30) captures the effect of inflation on the real value of short-term debt outstanding in period t, as well as its effect on the government surplus (accounting for the feedback from output to fiscal revenues). This margin generates only a temporary (within period) response of inflation. When debt is purely short-term, only contemporaneous inflation can adjust to ensure satisfaction of (29), since future inflation does not affect the real repayment of short-term liabilities. Moreover, the present value of the surplus component can be expressed as a function of inflation in period t.[23]

[21] Obviously for the canonical model we have $Rent_{t+j} = 0$ for all j.

[22] Intuitively, a more persistent inflation response would keep inflation elevated even when coupon payments have largely decayed. At that point, high inflation would no longer contribute meaningfully to debt sustainability. Conversely, if inflation were more front-loaded, then satisfying the intertemporal budget constraint would require a larger initial adjustment in inflation. Such a policy would fail to fully exploit the maturity structure of long-term debt; committing to a more persistent path allows inflation to be spread more efficiently over time, thereby reducing volatility. See Leeper and Zhou (2021) for further discussion of the role of maturity in shaping optimal inflation dynamics in these models.

[23] Specifically, in the log-linear model, the present value of the surplus can be written as $E_t \sum_{j\geq 0} \beta^j \Big(-\overline{G}\hat{G}_{t+j} + \widetilde{h}_Y(\hat{\pi}_{t+j} - \beta\hat{\pi}_{t+j+1}) + \widetilde{h}_\tau \hat{\tau}_{t+j} \Big)$ for some constants $\widetilde{h}_Y, \widetilde{h}_\tau$. Here, we used the Phillips curve to substitute output with inflation and taxes. The inflation terms simplify to $E_t \sum_{j\geq 0} \beta^j (\hat{\pi}_{t+j} - \beta\hat{\pi}_{t+j+1}) = \hat{\pi}_t$.

Note that, in a manner of speaking, equation (29) 'pins down' inflation in the canonical model. Following an unanticipated increase in government spending, the fiscal surplus falls in present value terms. Given that taxes do not fully offset the shock (and rents are zero), the planner sets inflation to ensure that the government's intertemporal budget constraint holds.

The logic underlying this argument parallels that of the Fiscal Theory of the Price Level. Although the surplus sequence is not exogenous to the policy problem, the resulting optimal policy equilibrium exhibits a similar feature: inflation adjusts so as to guarantee satisfaction of the intertemporal constraint.

In the extreme case in which all government debt is short-term, the planner effectively has no intertemporal margin through which to smooth the adjustment. Since future inflation does not affect the real value of short-term liabilities, the required adjustment must occur contemporaneously. As a result, current inflation is uniquely determined by the requirement that the budget constraint be satisfied. This is consistent with (30) when $\overline{b}_L = 0$.[24]

The Online Appendix E formalizes this property for a broader class of environments, including cases in which preferences are not quasilinear and government debt is not exclusively short-term. We adopt the following approach to demonstrate that the logic of the Fiscal Theory applies in the context of the canonical Ramsey model. We show that—generically—the optimal inflation policy in this model can be implemented through simple interest rate rules that set the nominal interest rate as a function of current inflation. These rules can be characterized analytically and are passive in the sense of Leeper (1991). From this finding, we conclude the relevance of the Fiscal Theory argument. For brevity, we refer the reader to the appendix E for a more detailed discussion.

Inflation under binding financial frictions. In contrast, the FTPL logic does not apply to the determination of optimal inflation under binding financial frictions. The key reason is that the left-hand side of (29) is itself a function of inflation (and of the debt sequence). Consequently, this equation is not the only relevant constraint to determine the equilibrium path of inflation. In the case in which debt is purely short-term, future inflation can adjust to ensure that (29) holds even without any contemporaneous response of inflation in period t.

Instead, in the presence of frictions, the optimal inflation path is jointly determined by (29) and the optimal sequence of short-term debt. This is evident from (26), which shows that inflation depends on the supply of short-term bonds $b_{S,t-1}$. Proposition 4 follows from combining (26) with the bond optimality condition (27).

The underlying logic parallels that in models such as Hagedorn (2018). When the real interest

[24]To further clarify, the claim we make here is not that (29) trivially determines a unique inflation path. This condition determines the cumulative response of prices/inflation to a spending shock, given the optimal tax sequence. The planner (effectively) selects the optimal inflation sequence to satisfy (29) which becomes the only relevant constraint, once one accounts for the Phillips curve induced output feedback effects on the tax base and real interest rates (standard in the FTPL). Ultimately, the optimal inflation sequence that satisfies (29) is determined by the maturity structure of debt, is the sense stated above. The tax sequence is matters only for determining the magnitude of the initial reaction of inflation (through the response of the present value of the surplus). The problem becomes trivial only in the case where debt is entirely short-term.

rate depends on the supply of government debt, the intertemporal budget constraint is not sufficient to determine the equilibrium price level sequence. Instead, the price sequence is jointly determined by monetary and fiscal/debt policies. The mechanisms emphasized in Hagedorn (2018) apply to our optimal policy problem.

In light of these remarks, we can interpret the components of inflation in equation (28). The first and third terms capture the effects of inflation on the real value of short- and long-term debt, which - as in the canonical model - generate a temporary and a persistent component of inflation. The loadings μ_1 and μ_3 however differ from those in the frictionless benchmark because these terms also reflect the endogenous response of short-term debt to the spending shock.[25]

The time-varying inflation-target component is driven by the rent term on the left-hand side of (29). A commitment to permanently higher inflation following a spending shock raises rents and thereby relaxes the intertemporal debt constraint. Importantly, the overall effect captured in equation (28) reflects not only the direct impact of inflation on rents, but also the endogenous response of short-term debt supply to the shock. We later characterize this response analytically.

Time-varying inflation target in the baseline model. Figure 2 confirms that the above properties hold in the baseline calibration of the model. The solid blue line corresponds to the case with binding frictions, while the dashed red line corresponds to the canonical model. As shown in the figure, inflation exhibits similar dynamics in both cases: a strong reaction on impact (short debt + surplus) followed by a transition through the long debt channel. However, whereas inflation in the canonical Ramsey model reverts to its steady-state value, in the model with frictions it stabilizes at a higher level. Optimal policy in the latter model implements a time-varying target.

[25] See appendix C for explicit formulae.

Figure 2: **Optimal Inflation**

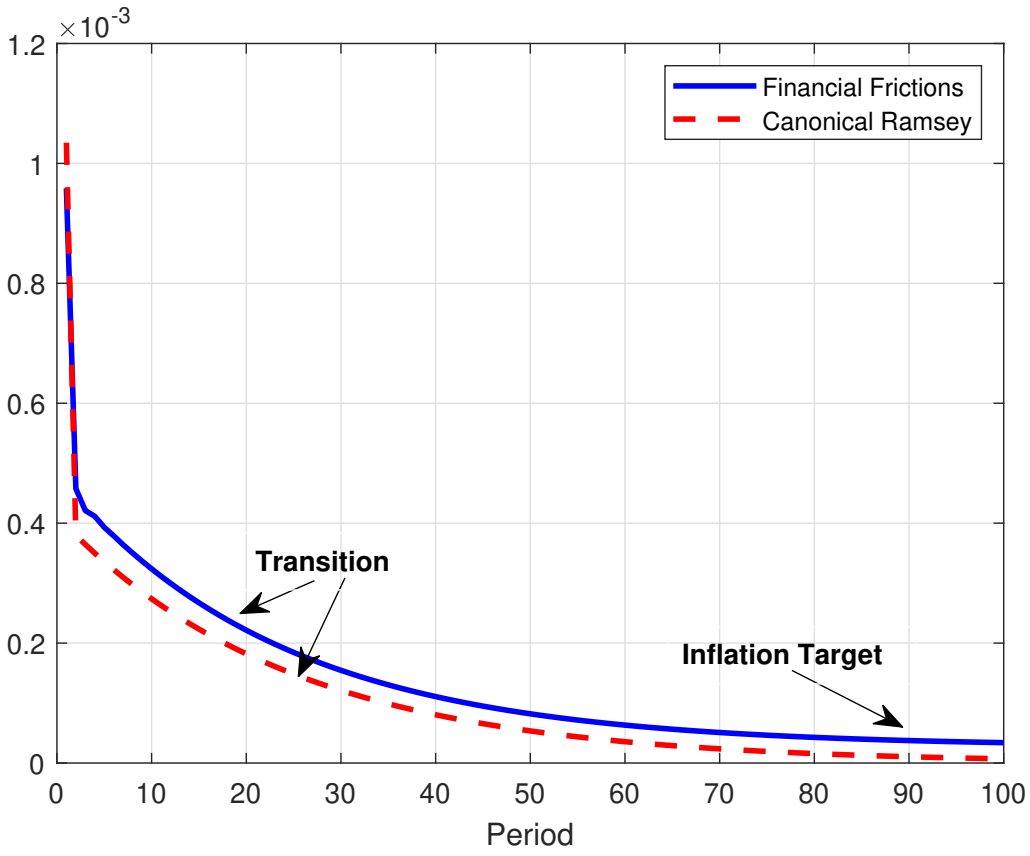

3.2 Optimal Taxes

3.2.1 Intertemporal tax distortions

We now turn to optimal taxation. We highlight the key differences in fiscal policy between the environment with financial frictions and the canonical benchmark. We focus on how the planner distorts taxes intertemporally to manipulate the real financing costs of debt — the margin where the two environments diverge most sharply. To isolate the effects on the optimal tax schedule, we restrict attention to the limiting case $\omega \to \infty$, thereby abstracting from inflation.

Analytical Tax Schedule. In Appendix B, we derive the tax schedule analytically and show that it can be expressed as a function of three components:

$$\tau_t = \mathcal{F}\Big(\text{Random-Walk Component},\ \text{Markup Management},\ \text{Intertemporal Distortions}\Big),$$

The first two components are easy to interpret in light of well-known results in the literature. The random walk property is the classical tax-smoothing result (Aiyagari et al., 2002). The Markup Management term is the endogenous response of the tax schedule to the wedge induced by monopolistic competition (see, e.g. Schmitt-Grohé and Uribe, 2004b). The government uses taxes to partially offset this markup-induced distortion, leading to a lower average tax rate in steady state, as well as lower taxes in high-spending states.

Our main focus in this section is on the third component of the tax schedule. To clarify how this influences taxes, we use explicit analytical expressions. Intertemporal distortions can be summarized as the following three elements:

$$D_{1,t} = \begin{cases} f_{\widetilde{\theta}_t}\left(\beta E_t \frac{u'(C_{t+1})}{u'(C_t)} - 1\right) & \textbf{FF} \\ 0 & \textbf{CM} \end{cases} \qquad D_{2,t} = \begin{cases} b_{t-1,S}u'(C_t)\left(\psi_{gov,t} - F_{\widetilde{\theta}_{t-1}}\psi_{gov,t-1}\right) & \textbf{FF} \\ b_{t-1,S}u'(C_t)\left(\psi_{gov,t} - \psi_{gov,t-1}\right) & \textbf{CM} \end{cases}$$

$$D_{3,t} = u'(C_t)\sum_{j\geq 0}\delta^j\left(\psi_{gov,t-j} - \psi_{gov,t-j-1}\right)b_{L,t-j-1}$$

where the abbreviations FF and CM denote the Ramsey model with financial frictions and the canonical model, respectively. Evidently, D_3 represents a common intertemporal distortion in both frameworks. The derivatives of the tax schedule with respect to the three components satisfy $\frac{d\tau_t}{dD_{1,t}} > 0$ $\frac{d\tau_t}{dD_{2,t}} < 0$ and $\frac{d\tau_t}{dD_{3,t}} < 0$.

We next examine each of the terms D_1, D_2, D_3. Our discussion complements the earlier interpretation of the first-order conditions of the Ramsey program, since D_1, D_2, D_3 appeared in the FONC for C_t. We thus analyze how the planner sets taxes to influence the consumption sequence and, through it, the real discount rates and the interest rate costs of issuing debt.

Real Discount Rates. Consider first $D_{2,t}$, focusing on an environment without frictions. To analyze this term, we study a one-off increase in spending. Because the shock tightens the implementability constraint, it permanently raises the value of the Lagrange multiplier.

The expression for $D_{2,t}$ implies that the planner will optimally target a lower tax rate in period t, while committing to permanently higher taxes in future periods (see Lucas and Stokey (1983); Faraglia et al. (2013), among others).

The intuition behind this well-known property becomes clear from a simple inspection of the intertemporal budget constraint in this frictionless model. Assuming, for convenience, that optimal long-term debt is zero in all periods, we have: $E_t \sum_{j\geq 0}\beta^j \frac{u'(C_{t+j})}{u'(C_t)}\text{Surplus}_{t+j} = b_{S,t-1}$.

A government that issues debt ($b_{S,t-1} > 0$) and commits to a sequence of positive future surpluses can increase the present value of its revenues by timing taxes to lower the marginal utility of consumption in period t (reducing the tax rate) while raising the marginal utility in subsequent periods (committing to higher taxes after t). In doing so, it effectively lowers the real discount factors for future revenues, $\beta^j \frac{u'(C_{t+j})}{u'(C_t)}$. Since the planner discounts the future at rate β, the optimal increase in taxes after period t is permanent.[26]

Consider now the case in which the short debt supply is insufficient to eliminate the frictions.

[26]Using the intertemporal constraint to build intuition for this policy is convenient as it enables us to summarize intertemporal prices in one equation; however, another way of interpreting the policy (common in many papers) is through the per-period budget constraint: $\beta E_t \frac{u'(C_{t+1})}{u'(C_t)} b_{S,t} + \text{Surplus}_t = b_{S,t-1}$. Lowering taxes in period t and raising them in $t+1$ increases the issuance price of debt. Beyond date t it is not optimal to distort prices as tax smoothing (the random walk property) suggests that $E_t D_{2,t+1} = E_t D_{2,t+2} = ... = 0$.

The term D_2 continues to reflect the same policy margin; however, unlike in the canonical model, it may continue to produce intertemporal tax distortions after $t+1$.

To see this, suppose that following the shock at t, the Lagrange multiplier remains permanently at a higher level $\psi_{gov,t}$ (i.e. there is no shock after t). We then have

$$D_{2,t+j} = \begin{cases} b_{t+j-1,S}\, u'(C_{t+j})\, \psi_{gov,t}\big(1 - F_{\widetilde{\theta}_{t+j-1}}\big), & j \geq 1, \\ \widetilde{\theta}_{\infty}\, \psi_{gov,t}\big(1 - F_{\widetilde{\theta}_{\infty}}\big), & j \to \infty, \end{cases}$$

Thus, generically, $D_{2,t+j} \neq 0$; moreover, the sequence $D_{2,t+j}$ may vary over time depending on the transitional dynamics of short-term government debt, private sector consumption, and—crucially—their ratio $\widetilde{\theta}$.

To examine the underlying mechanisms, it is useful to consider the intertemporal solvency condition. Maintaining that long bonds are in zero net supply (and that the discounted present-value series converges) we have:

$$E_t \sum_{j\geq 0} \beta^j \frac{1}{u'(C_t)} \Pi_{k=0}^{j-1} F_{\widetilde{\theta}_{t+k}} \left[\int_{\widetilde{\theta}_{t+j}}^{\infty} \theta dF_\theta + u'(C_{t+j}) \text{Surplus}_{t+j} \right] = b_{t-1,S}$$

A policy that lowers taxes in t and promises to raise them in $t+1$ again affects the real discount rate. However, unlike the frictionless scenario, the discount factor governing this intertemporal tradeoff is now $\beta \frac{u'(C_{t+1})}{u'(C_t)} F_{\widetilde{\theta}_t}$ (rather than $\beta \frac{u'(C_{t+1})}{u'(C_t)}$). Analogously, for the intertemporal tradeoff of taxes between t and $t+j$ the discount factor is $\beta^j \frac{u'(C_{t+j})}{u'(C_t)} \Pi_{k=0}^{j-1} F_{\widetilde{\theta}_{t+k}}$. Thus, the effective discount factor applied to future surpluses differs from the social discount factor. As long as the economy remains in the constrained regime and $\widetilde{\theta}$ exhibits transitional dynamics, the intertemporal tradeoffs faced by the planner vary over time and taxes need not remain constant after period t.

Exactly how the optimal tax schedule evolves from $t+1$ onward is difficult to derive analytically. A numerical example will illustrate that $D_{2,t}$ behaves quite differently in the model with frictions than in the canonical model, showing that it has a persistent effect on the tax schedule.

Liquidity Cutoff. Next, consider the term $D_{1,t}$. This generates an intertemporal distortion in the tax schedule, which arises solely due to financial frictions. This term captures the indirect impact of taxes on consumption and through that on the ratio $\widetilde{\theta}_t$ and the liquidity premium of short-term debt.[27] Recall that this term corresponds to the "liquidity cutoff" appearing in (18). A marginal decrease in $\widetilde{\theta}_t$ reallocates a mass $f_{\widetilde{\theta}_t}$ of agents from the constrained region to the unconstrained region. For these marginal agents, the net benefit from holding short-term government debt is given

[27] In this real model we have $K_t \equiv u'(C_t) p_{t,S} b_{S,t} = \int_{\widetilde{\theta}_t}^{\infty} \theta dF_\theta + b_{S,t} \beta F_{\widetilde{\theta}_t} E_t u'(C_{t+1})$ which measures the utility cost associated with the issuance of short-term government debt. Differentiating with respect to $\widetilde{\theta}_t$ yields

$$K_t' = -f_{\widetilde{\theta}_t} \widetilde{\theta}_t + b_{S,t} \beta f_{\widetilde{\theta}_t} E_t u'(C_{t+1}) = \widetilde{\theta}_t f_{\widetilde{\theta}_t} (\beta E_t \frac{u'(C_{t+1})}{u'(C_t)} - 1).$$

by the wedge, $\beta E_t \frac{u'(C_{t+1})}{u'(C_t)} - 1$. Consequently, changes in the cutoff affect the value of government debt issuance in proportion to this wedge.

To see how $D_{1,t}$ induces intertemporal tax distortions, consider again an unanticipated increase in spending. Following the shock, $\beta E_t \frac{u'(C_{t+1})}{u'(C_t)}$ will drop (consumption falls initially and gradually rises toward a new steady state). Responding to $D_{1,t}$, the planner will lower taxes in t and gradually increase them, as the economy transitions to the steady state. The objective is to partially offset the resulting gap in $\beta E_t \frac{u'(C_{t+1})}{u'(C_t)} - 1$ and thus to mitigate the rise in the interest rate cost of short-term debt.

This mechanism is not confined to periods t and $t+1$. It applies to the entire transition path from the moment the economy experiences the shock and until it reaches the new long run equilibrium, as shown below.

Long bond price. Finally, the term $D_{3,t}$ is analogous to $D_{2,t}$ however it pertains to long-term debt. When the supply of long bonds is positive, the intertemporal constraint is

$$E_t \sum_{j\geq 0} \beta^j \frac{1}{u'(C_t)} \Bigg[\underbrace{\int_{\widetilde{\theta}_{t+j}}^{\infty} \theta dF_\theta + (F_{\widetilde{\theta}_{t+j}} - 1) u'(C_{t+j+1}) b_{S,t+j} \beta}_{\text{Liquidity/Safety rents}} + u'(C_{t+j}) \text{Surplus}_{t+j} \Bigg]$$
$$= b_{t-1,S} + b_{t-1,L} E_t \sum_{j\geq 0} \delta^j \beta^j \frac{u'(C_{t+j})}{u'(C_t)}$$

The timing of taxes now impacts both the real discount rates of future surpluses, and the real price of outstanding long-term debt on the right-hand side of the constraint.

An increase in government spending results in progressively higher taxes. According to our formula, $\frac{D_{3,t+j}}{u'(C_{t+j})} = \delta^j \Delta\psi_{gov,t} b_{L,t-j-1}$ and therefore, the distortion created loads geometrically on future dates governed by the duration parameter δ. Higher δ leads to a smoother tax sequence. In the limit, $\delta = 1$, $\frac{D_{3,t+j}}{u'(C_{t+j})}$ follows a random walk which implies no intertemporal tax distortions. These properties are well known in the literature (see, e.g., Lucas and Stokey, 1983; Leeper and Zhou, 2021).

Numerical Example. To illustrate how the above mechanisms translate into intertemporal tax dynamics, we present a numerical example in Figure 3. The top panels plot the responses of taxes and D to a positive spending shock in the model with financial frictions. The bottom panels show the analogous responses in the canonical model.[28]

Consider first the right panels, which focus on the terms D. The red dashed line traces the path of $D_{1,t}$ following the shock. Upon impact, this term turns positive and then gradually converges to the (new) steady-state value.

The left panel shows how this component affects the tax schedule. The solid/blue line depicts the equilibrium path of the tax rate. The dashed/red line subtracts the contribution of $D_{1,t}$ from

[28] As in the previous section, we compute a log-linear solution of the models. This allows us to fix the quantities $\overline{b}_S, \overline{b}_L$ in the frictionless economy at the steady-state levels obtained under financial frictions. We maintain fully rigid prices in these simulations.

the response, showing how taxes would evolve if the planner ignored this margin. $D_{1,t}$ lowers taxes and generates an intertemporal distortion, as the planner gradually increases the tax rate over time to reduce the real interest rate on short-term debt. These effects are absent in the canonical model since $D_{1,t} = 0$ in this case.

Figure 3: **Taxes and Intertemporal Distortions**

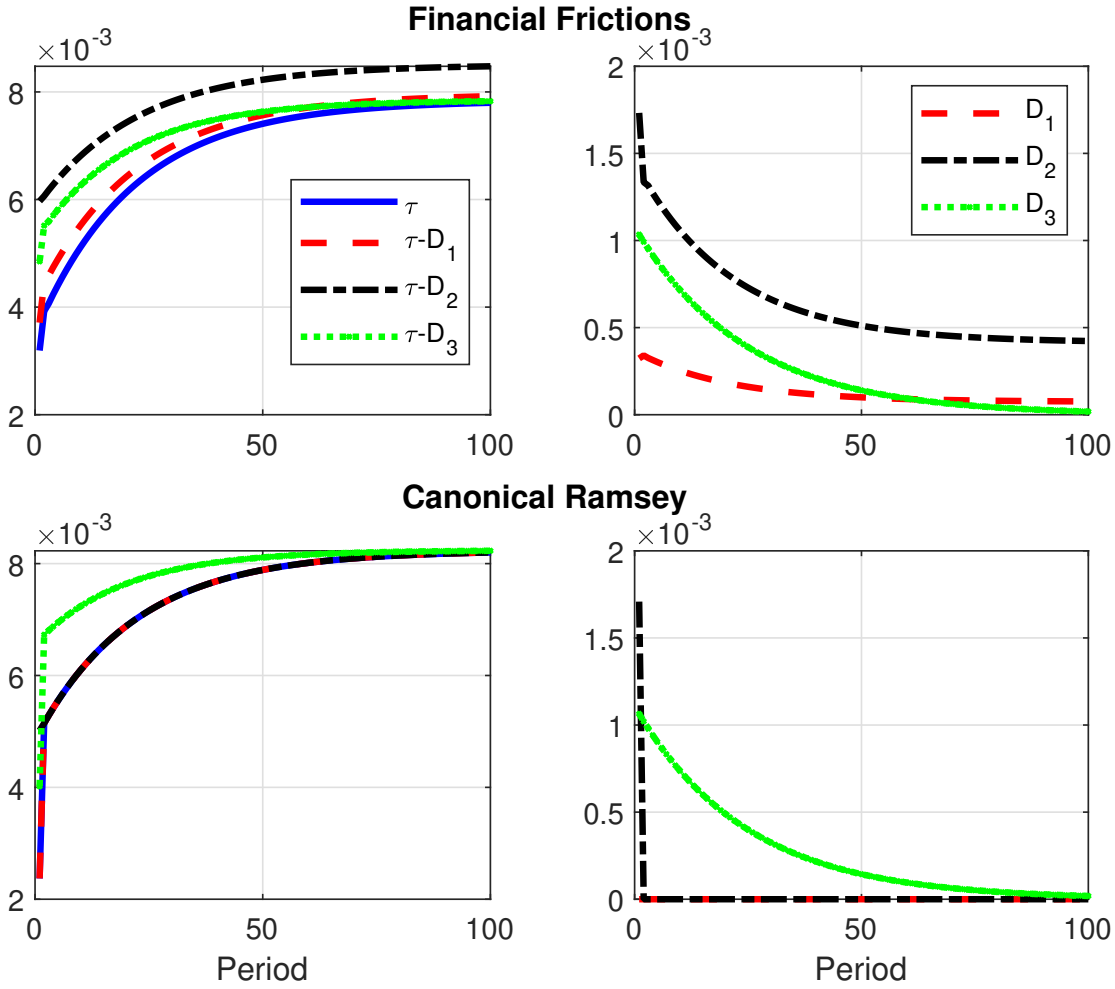

Notes: The figure plots the response of taxes to a spending shock in the log-linear solution of the model, decomposing the underlying sources of intertemporal tax distortions (see text for details).

Next, consider the role of the term $D_{2,t}$, depicted by the black dashed line in the right-hand panels. In the canonical model, this component generates a distortion only on impact. By contrast, in the presence of financial frictions, $D_{2,t}$ displays a pronounced initial response to the shock and then converges gradually to a strictly positive long-run steady-state value. This dynamic path reflects both the transitional adjustment of short-term debt and consumption, and the eventual permanent decline in the ratio $\widetilde{\theta}$.

The top-left panel highlights the quantitative importance of $D_{2,t}$ for the tax schedule. Without this channel, taxes rise substantially on impact and follow a flatter transition path.

Finally, the green lines correspond to the term $D_{3,t}$. Unsurprisingly, its effect on the tax schedule is essentially identical in the models with and without frictions.

3.3 Inflation and tax volatility

The previous paragraphs highlighted two central findings regarding optimal inflation and tax policy in the presence of financial frictions. First, inflation follows a process that contains a random-walk component. Second, optimal tax policy operates through novel intertemporal distortion margins, which generate persistent effects on the tax schedule.

A large literature on optimal Ramsey allocations finds that taxes bear most of the burden of adjustment following a fiscal shock, while optimal policy assigns a very high weight to stabilizing inflation (see, for example, Schmitt-Grohé and Uribe, 2004b; Siu, 2004; Lustig et al., 2008; Faraglia et al., 2013, among others). In Appendix F, we use the baseline model with sticky prices to examine whether financial frictions—and the resulting differences in optimal monetary and fiscal policy identified above—alter this conclusion.

Our results are as follows: though inflation variability is larger relative to the canonical model and taxes are less volatile, the differences are not considerable. To put it differently, under financial frictions, taxes continue to absorb most of the adjustment following spending shocks. Moreover, inflation persistence is higher than in the canonical model, with this effect becoming more pronounced when the model is approximated around an equilibrium with low levels of long-term debt.

These results are perhaps surprising given our findings in subsection 2.3.2, where the same baseline calibration implied an optimal steady-state inflation target of 2 percent together with lower taxes in the equilibrium with frictions. Thus, although inflation fluctuations over the business cycle are modest, the planner nevertheless chooses to generate sizable steady-state inflation distortions. It is as if the planner faces a different trade-off between inflation and taxes in steady state than over the business cycle. A more detailed investigation of this implication of the model, though interesting, is left for future work.

4 Optimal Portfolios

This section presents our main analysis of the optimal management of government portfolios. We begin by establishing connections to the related literature. We then describe the numerical algorithm used to solve the model, which accounts for the presence of multiple equilibria. Finally, we present and discuss our results.

4.1 Effectively Complete vs Incomplete Markets

We briefly discuss a benchmark result from the existing debt management literature. In particular, we revisit the finding of Angeletos (2002) and Buera and Nicolini (2004) (henceforth ABN) on optimal debt policy under complete markets. We then explain how our optimal policy program with financial frictions departs from this benchmark, and clarify the role of short-term debt in our environment. Although the ABN result does not directly apply in our setting, it remains informative

both conceptually and for the numerical algorithm we employ.

Throughout this subsection, we focus on an economy with perfectly rigid prices. This suffices to illustrate the key mechanisms.

First, to introduce notation, let $z_t(\mathbf{g}^t)$ denote the expected present discounted value of government surpluses and liquidity/safety rents, where $\mathbf{g}^t \equiv (G_0, G_1, \ldots, G_t)$. Let $\Upsilon_t(\mathbf{g}^{t-l})$, for $l \in \{0, 1\}$, denote the market value of government debt. The intertemporal budget constraint can be written compactly as $\Upsilon_t(\mathbf{g}^{t-l}) = z_t(\mathbf{g}^t)$. When $l = 0$, the market value is measurable with respect to the contemporaneous history $\mathbf{g}^t$; when $l = 1$, it is measurable only with respect to $\mathbf{g}^{t-1}$.

We refer to an equilibrium as being under *effectively complete markets* when $\Upsilon_t(\mathbf{g}^t)$ can be chosen residually such that $\Upsilon_t(\mathbf{g}^t) = z_t(\mathbf{g}^t)$ almost surely for all $t \geq 1$. In this case, the date-0 condition $\Upsilon_0(\mathbf{g}^0) = z_0(\mathbf{g}^0)$ is sufficient for an equilibrium. By contrast, under incomplete markets the date-0 implementability condition is not sufficient, and the sequence $\Upsilon_t(\mathbf{g}^t) = z_t(\mathbf{g}^t)$ for $t \geq 1$ is binding in the Ramsey problem.[29]

Consider first the optimal policy equilibrium in the canonical model. Suppose spending follows a two-state Markov process taking values $G^+ > G^-$. Standard results imply that under complete markets z_t converges to a two-state process satisfying $z^+ < z^-$ and consumption converges to $C^+ < C^-$.[30]

Given this allocation, satisfying the intertemporal constraints requires choosing bond holdings such that $b_{S,t-1} + b_{L,t-1}\big(1 + \delta q_{t,L}(\mathbf{g}^t)\big) = z_t(\mathbf{g}^t)$ for all t, where $q_{t,L}$, the price of long-term debt, is itself a two-state process with $q_L^+ < q_L^-$. With two bonds and two states, this condition defines a square linear system in (b_S, b_L).

In Appendix B we provide analytical examples illustrating the following well-known result.

Result (ABN). *The optimal bond portfolio in the canonical model satisfies $b_S < 0$ and $b_L > 0$.*

Therefore, the government optimally saves in short-term assets and issues long-term debt. The intuition is that long-term bonds provide *fiscal hedging*, insuring the government budget against spending shocks through the negative covariance between bond prices and fiscal deficits. By issuing long-term debt, the planner ensures that when spending rises, the market value of outstanding debt is reduced sufficiently so that $\Upsilon_t(\mathbf{g}^t)$ replicates the state-contingent variation in z_t.

The ABN result relies on two key conditions. First, short- and long-term bond positions must be free-choice variables. Second, the government spending process must have the same finite dimensionality as the available bond structure.

Neither condition holds in the program we solve in this section. We do not restrict government spending to a two-state Markov chain; instead, we use empirically plausible calibrations in which spending follows a continuous distribution. With two bonds but infinitely many states, the asset structure cannot span the state space, and exact market completion is impossible.

[29] See Angeletos (2002); Buera and Nicolini (2004); Faraglia et al. (2019) among others.

[30] This result may hold asymptotically. Convergence in finite time requires additional assumptions, for example that long-term bonds are consols. See appendix B for further details.

More importantly, financial frictions restrict the set of feasible short-term positions. Since $b_S < 0$ is not admissible given the specification of household preferences, the government cannot accumulate short-term assets as in the ABN benchmark. More generally, short-term debt is not a free variable: liquidity/safety provision constrains its level in equilibrium, so $\Upsilon_t(\mathbf{g}^t)$ cannot be adjusted arbitrarily to match $z_t(\mathbf{g}^t)$.[31] Consistent with our formulation of the Ramsey policy program in previous sections, the constraints for all $t \geq 1$ must be explicitly accounted for in the Ramsey policy.

Despite these differences, the benchmark ABN result remains informative in our context. Recall that, in our model, there exist regions of the state space in which the policy problem is effectively equivalent to the canonical frictionless benchmark. In particular, when the supply of short-term debt is sufficiently high, the economy features abundant liquidity and safety, and the distortions induced by frictions vanish. In such cases, the logic underlying ABN re-emerges: although short-term debt remains strictly positive, the planner has an incentive to tilt the portfolio toward the ABN benchmark to capture the hedging benefits of long-term debt.

This insight informs the construction of our numerical algorithm. Specifically, in regions of abundant liquidity, we impose the ABN allocation when computing the optimal portfolio; that is, we restrict attention to allocations that maximize long-term debt issuance subject to remaining within the high-liquidity region. In the next subsection, we explain why this is a necessary and justified restriction in our setting.

Short-term debt as a shock absorber. Our model differs from the canonical framework in one additional important respect. In our framework, the right-hand side of the constraint, $b_{S,t-1} + b_{L,t-1}\big(1 + \delta q_{t,L}(\mathbf{g}^t)\big) = z_t(\mathbf{g}^t)$ is also a function of government debt. Consequently, $z_t(\mathbf{g}^t)$ cannot be treated as exogenous to the debt management problem. We now clarify the role of short-term bonds in satisfying the government's intertemporal constraints. To isolate the relevant mechanism, we consider again the simplest version of the model, assuming quasilinear preferences. In this case, and under rigid prices, $q_{t,L}$ is constant. The implementability condition therefore becomes $\Upsilon_t(\mathbf{g}^{t-1}) = z_t(\mathbf{g}^t)$. In Appendix C, we prove the following result:

Proposition 5 *Consider the optimal policy under quasilinear utility and stochastic government spending. An unanticipated increase in spending leads to:*

(i) a permanent increase in the tax rate;

(ii) a decline in short-term debt $b_{S,t}$ and an increase in long-term debt $b_{L,t}$;

(iii) an increase in total government debt.

Thus, the planner will lower the supply of short-term bonds in periods of high spending needs. Intuitively, a lower $b_{S,t}$, which is 'financed' through a higher issuance of long-term bonds enables the planner to increase the rent implicit in z_t and thus mitigate the increase in taxes required to finance spending. Short-term bonds act as a *shock absorber*.

[31] In Appendix B, we provide numerical examples illustrating this property. We solve the Ramsey policy under the date 0 implementability constraint and illustrate that $\Upsilon_t(\mathbf{g}^t) = z_t(\mathbf{g}^t)$, $t \geq 1$ cannot hold.

The result in Proposition 5 continues to hold in the baseline model with sticky prices and curved utility. Unlike the quasilinear model however, in this case long-term government debt provides fiscal insurance (and therefore also acts as a shock absorber) through the endogenous price fluctuation channel we previously mentioned.

4.2 Numerical Algorithm

4.2.1 State Variables

In this section, we describe our numerical solution method. We first define the state vector X. Marcet and Marimon (2019) develop a recursive formulation of dynamic optimization problems with forward-looking constraints by augmenting the state space with appropriate co-state variables (Lagrange multipliers). In our model with long-term debt featuring coupon payments that extend over an infinite horizon, the state vector is, in principle, infinite-dimensional. The planner's commitments regarding inflation and consumption extend indefinitely into the future, implying that an infinite number of lags of multipliers, bond quantities, and inflation would need to be tracked as state variables.

A closer inspection of the optimality conditions derived in previous sections, however, reveals that a much lower-dimensional representation of the state vector is possible. In particular, the following recursive state variable summarizes the planner's policy commitments: $\Psi_t = \frac{b_{L,t-1}\Delta\psi_{gov,t}}{\pi_t} + \frac{\delta}{\pi_t}\Psi_{t-1}$ [32]

In the Online Appendix D we apply the arguments of Marcet and Marimon (2019) to derive the saddle-point functional equation (SPFE) and the associated welfare objective function. The welfare objective is $W(X_t) = W(b_{S,t-1},\, b_{L,t-1},\, \widetilde{\theta}_{t-1},\, \psi_{gov,t-1},\, \psi_{PC,t-1},\, \Psi_{t-1},\, G_t)$.

The dimensionality of the state vector can be further reduced by eliminating $\psi_{PC,t-1}$, since this variable is functionally determined by $\psi_{gov,t-1}$.[33] The resulting state vector is therefore

$$X_t \equiv \left(b_{S,t-1},\, b_{L,t-1},\, \widetilde{\theta}_{t-1},\, \psi_{gov,t-1},\, \Psi_{t-1},\, G_t\right).$$

4.2.2 Parameterizing Expectations

Solving the SPFE explicitly is difficult. The value function W depends on 6 state variables, including Ψ, whose value accumulates over time and may therefore vary over a wide range along the equilib-

[32] Iterating backward yields

$$\Psi_t = \sum_{j\geq 0} b_{L,t-1-j}\frac{\Delta\psi_{gov,t-j}}{\prod_{k=0}^{j}\pi_{t-k}}.$$

This term corresponds to the discounted sum of past multipliers associated with the government budget constraint that appears in the first-order conditions.

[33] We focus on the solution to the optimal policy problem initiated with multipliers $\psi_{gov,-1} = \psi_{PC,-1} = \Psi_{-1} = 0$. Along the optimal path,

$$\psi_{PC,t} = -\frac{\omega}{\eta}\psi_{gov,t}$$

holds for all t, so $\psi_{PC,t}$ is fully determined by $\psi_{gov,t}$ and can be omitted from the state vector.

rium path. As a result, applying standard value function methods on a discretized state space is particularly difficult in our setting.

We solve the model using the system of first-order conditions. We apply the PEA (Den Haan and Marcet, 1990) to numerically solve the system of equations comprising the first-order conditions (10)–(16), together with the constraints of the planner's problem.

The PEA approximates the conditional expectations appearing in these equations with polynomials of the state variables. Our approximations are made with polynomials

$$\left\{\Phi_{(q,S)}(X_t,\kappa_{q,S}),\Phi_{(q,L)}(X_t,\kappa_{q,L}),\Phi_{\psi,S}(X_t,\kappa_{\psi,S}),\Phi_{\psi,L}(X_t,\kappa_{\psi,S}),\Phi_{PC}(X_t,\kappa_{PC})\right\}$$

where the subscripts indicate which of the expectations the function $\Phi_{()}$ approximates. Specifically, (q,S) $((q,L))$ is used for the conditional expectation appearing in the short (long) bond price equation and (ψ,S) (ψ,L) for the conditional expectations appearing in the FONCs for short (long) bonds (the RHS variables in (15) and (16) respectively). (PC) is the approximation of the expected inflation term in the Phillips curve. Finally, $\kappa_{(.)}$ are the coefficients of the variables in the state vector (their squares, cubes, cross products depending on the order of the approximation being used).[34]

The PEA initiates the approximations $\Phi_{()}$ and, solving the model iteratively (using simulated samples), it finds a fixed point for the coefficients $\kappa_{(.)}$. Given that this procedure is well known, we refer the reader to the appendix D for a more complete description. In what follows, we discuss how we adapt the PEA to deal with the multiplicity of solutions to the system of first-order conditions.

4.2.3 Multiplicity/ Debt limits

Central to our approach to addressing the multiplicity is the introduction of debt constraints on bond positions. We begin by defining these objects and explaining how they help us recover the global optimum. Following Aiyagari et al. (2002) and Faraglia et al. (2016, 2019), we assume that the issuance of short- and long-term bonds is subject to ad hoc limits of the form

$$b_{x,t}\in[\underline{M}_x,\overline{M}_x],\qquad x\in\{S,L\},\tag{31}$$

where $|\underline{M}_x|,|\overline{M}_x|<\infty$.[35]

The prior literature has imposed similar debt limits for two main reasons. First, from a numerical perspective, imposing ad hoc limits improves the stability and performance of the solution algorithm (Maliar and Maliar, 2003; Faraglia et al., 2019, 2016; Judd, Maliar, and Maliar, 2011). Second, these bounds also admit a meaningful economic interpretation. In particular, the upper bound $\overline{M}_L$ can be interpreted as a constraint that rules out government overborrowing—situations in which public debt becomes sufficiently large that the economy operates on the wrong side of the Laffer curve.

[34] See appendix D for the exact specification of these approximations.

[35] To account for these debt limits we augment the program with inequality constraints. See Appendix D for the resulting FONCs and a discussion of our approach to solve the model with occasionally binding constraints.

Analogously, lower bounds on asset positions have been motivated in the related literature either as a way to rule out overborrowing by the private sector or to preclude negative government asset positions altogether (Lustig et al., 2008; Faraglia et al., 2019).[36]

Long-term debt limits. In our numerical experiments below, we assume a loose upper bound for long-term debt. Specifically, we set $\overline{M}_L = 0.3$, which implies a maximum quantity of long-term bonds equal to 30 percent of steady-state GDP. In terms of market value (bond quantities multiplied by prices), this allows long-term debt to reach approximately 167 percent of annual steady-state GDP.

For the lower bound, we consider alternative scenarios. Our baseline model sets $\underline{M}_L = 0$, ruling out government holdings of long-term private debt (as in Faraglia et al., 2019; Lustig et al., 2008; Nosbusch, 2008).[37] We also consider alternative specifications that allow for government lending to the private sector ($\underline{M}_L < 0$). We return to these versions of the model in the final paragraph of this section.

Short-term debt limits. Short-term debt issuance is not subject to a lower bound constraint, as this would not be binding in our model. We can thus interpret $\underline{M}_S = 0$ as the natural debt limit for households.

To determine the upper bound $\overline{M}_S$, we proceed as follows. Given our calibration of the density function $F_{\widetilde{\theta}}$ (see below), we identify the interval of the short-term bond quantity over which $F_{\widetilde{\theta}} \approx 1$, or more precisely, $F_{\widetilde{\theta}} \geq 1 - \epsilon$, where ϵ is a small positive number. We then set

$$\overline{M}_S = \sup\left\{ b_S : F_{\widetilde{\theta}} < 1 - \epsilon \right\} \tag{32}$$

as the upper debt limit. In our simulations $\overline{M}_S$ is a constant real number.

This choice requires a brief explanation. Recall that in models with perfect foresight, the debt composition is not determined when $F_{\widetilde{\theta}} = 1$. In our theoretical analysis, it was sufficient to verify the optimum at a single point $\overline{M}_S$, realizing that for $b_S > \overline{M}_S$ the objective welfare function is essentially flat.

We show that an analogous property holds in the stochastic equilibrium. To show that it practically holds in the PEA approximation of the equilibrium, consider the two first-order conditions for

[36] See Aiyagari et al. (2002) for a formal definition of the *natural borrowing limit* in this class of models.

[37] This is typically motivated by the observation that the consolidated government sector rarely holds private assets in normal times to avoid exposure to private default risk. Notable exceptions occurred during the 2008–09 financial crisis and again during the COVID-19 pandemic, when central banks purchased large quantities of private assets. We return to these observations later.

short- and long-term bonds when $F_{\widetilde{\theta}} \approx 1$ and assuming loose upper bound debt limits.

$$\psi_{gov,t} = \frac{E_t \psi_{gov,t+1} u'(C_{t+1})/\pi_{t+1}}{E_t u'(C_{t+1})/\pi_{t+1}} = \frac{\Phi_{(\psi,S)}(X_t, \kappa_{(\psi,S)})}{\Phi_{(q,S)}(X_t, \kappa_{(q,S)})} \tag{33}$$

$$\psi_{gov,t} = \frac{E_t \psi_{gov,t+1} \sum_{j\geq 1} \beta^j \delta^{j-1} \frac{u'(C_{t+j})}{\Pi_{k=1}^{j} \pi_{t+k}}}{E_t \sum_{j\geq 1} \beta^j \delta^{j-1} \frac{u'(C_{t+j})}{\Pi_{k=1}^{j} \pi_{t+k}}} = \frac{\Phi_{(\psi,L)}(X_t, \kappa_{(\psi,L)})}{\Phi_{(q,L)}(X_t, \kappa_{(q,L)})} \tag{34}$$

Equations (33) and (34) cannot determine an optimal portfolio. Since X_t contains only predetermined endogenous variables (33) and (34) give us two equations to solve for $\psi_{gov,t}$. Therefore, the multiplier is overdetermined while the values for bond holdings are indeterminate. This indeterminacy is not present when $F_{\widetilde{\theta}} < 1$ as $b_{S,t}$ appears in the planner's first-order condition.[38]

By choosing $\overline{M}_S$ as in (32), we essentially rule out the possibility of indeterminate portfolios. The optimum may arise either as an interior solution or as a corner solution at $b_S = \overline{M}_S$.[39]

Implicitly, we are assuming that the planner's value function is (weakly) decreasing in b_S past point $\overline{M}_S$, so that substituting long-term debt for short-term government bonds (setting $b_{S,t} > \overline{M}_S$) is never optimal.

The previous paragraph provided the rationale for this assumption. In the canonical model, it is optimal for the government to issue long-term debt to benefit from fiscal insurance, financing this position with short-term savings. We expect a similar property to hold in the limit $F_{\widetilde{\theta}} = 1$ in our model and construct our algorithm accordingly.

Finally, given that the bound $\overline{M}_S$ that satisfies (32) is around 0.3 (approximately 30 percent of GDP) our assumption is not restrictive when compared with the canonical model. For example, Faraglia et al. (2019) solve incomplete market models with two assets and find that even if both assets are constrained to not be negative, the government never finds it optimal to issue large amounts of short-term debt. In their simulations, positive short-term bond issuance occurs when b_L is close to

[38] In this case, the short bond equation holds for a specific value $b_{S,t}$ and an optimal portfolio can be found.

[39] Notice that in practice it is not possible to determine the exact value of $\overline{M}_S$ corresponding to the corner solution (i.e., as in the analytical models of the previous sections). Since both b_S and C determine $\widetilde{\theta}$, the value of $\overline{M}_S$ satisfying (32) can only be determined either in steady state or approximately in stochastic simulations. As a result, for a fixed $\overline{M}_S$ in the stochastic model, there may be cases in which the first-order conditions are numerically satisfied for values of b_S close to $\overline{M}_S$ as well as exactly at $\overline{M}_S$. We can nevertheless rank all such critical points using the welfare function approximation discussed below.

The potential persistence of this multiplicity is discussed further in the Appendix D, where we also consider alternative approaches to address the multiplicity near the bound. For example, we experimented with introducing a "soft upper bound" on short-term debt by adding a transaction cost function $\mathcal{T}_S = \nu(b_{S,t} - \overline{M}_S)^2$ when $b_{S,t} > \overline{M}_S$, and zero otherwise. For a small cost parameter ν, the optimal policy program changes only minimally, but the indeterminacy problem is avoided. In particular, when $F_{\widetilde{\theta}} = 1$, the first-order condition for $b_{S,t}$ becomes

$$\beta \psi_{gov,t} E_t \frac{u'(C_{t+1})}{\pi_{t+1}} - \beta E_t \left[\psi_{gov,t+1} \frac{u'(C_{t+1})}{\pi_{t+1}} \right] - 2\nu(b_{S,t} - \overline{M}_S) u'(C_t) = 0.$$

Our results with the transaction cost are very similar to those in the baseline specification.

Finally, section D of the appendix also discusses the approach of Faraglia et al. (2019) to deal with the indeterminacy of optimal portfolios in the canonical model and explains the advantages of our approach in the model considered here.

its upper bound.[40] We choose a looser upper bound $\overline{M}_L$ than in that paper precisely to avoid cases in which the planner is forced to increase b_S because of the long-term debt issuance constraint.

SPFE approximation. Even with a tight upper bound on the short bond issuance, our program remains non-convex and the multiplicity of solutions remains an inherent feature of the system of first-order conditions. The first-order conditions may be simultaneously satisfied at an interior optimum $F_{\widetilde{\theta}} < 1$ and at the bound $\overline{M}_S$ (or approximately satisfied for points close to $\overline{M}_S$).

We complement the PEA with an approximation of the planner's objective function that solves the SPFE. Our approximation is:

$$W(X_t) \equiv E\left[\sum_{j\geq 0} \beta^j \widetilde{U}_{t+j} \mid X_t\right] \approx \Phi_W(X_t, \kappa_W) \tag{35}$$

where $\widetilde{U}_{t+j}$ has been previously defined as the household's utility function and $E(.|X_t)$ denotes the expectation conditional on the date t state variables.Using (35), we can evaluate and rank all critical points satisfying the first-order conditions, thereby identifying the welfare-maximizing solution. (See Appendix D for a detailed discussion).

4.3 Calibration

We now describe our calibration strategy. We follow Ghomi et al. (2026) in our choice of functional forms and parameter values. Importantly, parameters are not chosen to match the Ramsey steady state of the model. Instead, we calibrate an auxiliary model with ad hoc interest rate and tax rules (described in detail in Ghomi et al., 2026). The main target of the calibration is to match the empirical findings of Greenwood et al. (2015).

The model is calibrated at a quarterly frequency. We set $\beta = 0.995$, targeting an annual real return on long-term debt of approximately 2 percent.[41] We set the long-term bond decay parameter $\delta = 0.96$, implying an average maturity of long-term debt of about 25 quarters. Given a short-term-to-long-term debt ratio of 12.5 in U.S. data (see, e.g. Ghomi et al., 2026), this implies an overall debt maturity of about 5 years. We also set the steady-state debt-to-GDP ratio to 60 percent at an annual frequency.

Parameters ω and η are set to 17.5 and -6, respectively, following Schmitt-Grohé and Uribe (2004b). We set $\gamma = 1$, implying a unit Frisch elasticity of labor supply, a standard value in the macroeconomics literature. Finally, the persistence of the spending shock, ρ_G, is set to 0.95, and its standard deviation to 0.0075.

We next describe the choice of f_θ and the fiscal and monetary policy rules. First, we calibrate $\overline{q}_S$

[40] See the simulations of the buyback model under "no lending" in Figure 7 of Faraglia et al. (2019). Lustig et al. (2008); Nosbusch (2008) also solve Ramsey models with ad hoc borrowing constraints and find that optimal portfolios remain tilted toward long-term debt.

[41] The net inflation rate is zero in the deterministic steady state of the auxiliary model.

to match a 1-percentage-point annual term premium. Second, in the log-linear model, taxes follow the feedback rule $\hat{\tau}_t = \phi_\tau \hat{B}_{t-1}$, where $\hat{B}$ denotes the face value of total real debt (see, e.g., Leeper, 1991). Monetary policy follows an inertial Taylor rule $\hat{i}_t = \rho_i \hat{i}_{t-1} + (1-\rho_i)\phi_\pi \hat{\pi}_t$. We set $\phi_\tau = 0.01$, $\rho_i = 0.9$, and $\phi_\pi = 1.25$.[42]

Given these choices, our principle in calibrating the distribution f_θ is as follows. We assume that f_θ is log-normal, leaving two parameters (mean and variance) to match relevant targets. We calibrate the mean so that steady-state consumption equals 80% of output, which we normalize to one.

Finally, we set the variance to match the estimated elasticity of the term premium with respect to the short-term debt-to-GDP ratio in Greenwood et al. (2015). This paper finds that a one percent increase in the ratio reduces the annualized spread between T-bills and T-notes/bonds by about 16 basis points for four-week bills and about 8 basis points for ten-week yields. We target a 2 basis points change in the spread, corresponding to our quarterly model.[43]

4.4 Numerical Results

Short- and long-term debt. Figure 4 presents simulations of several endogenous variables in our baseline model. Specifically, we show 1,000 periods in which government spending follows the stochastic process defined earlier.[44]

The top-left panel traces the evolution of short-term debt in the sample; the top-right plots the quantity of long-term bonds. The bottom-left panel shows the (quarterly) market value of long-term debt, $q_{L,t} b_{L,t}$. Finally, the bottom-right panel displays the annualized inflation rate.

Consider first the top-left panel. The red horizontal line denotes the upper debt limit, $\overline{M}_S$. The graph shows that short-term debt fluctuates well below this upper bound, essentially always remaining in the region where financial frictions bind.

Our previous theoretical results help interpret this prediction of the model. In Section 2.3, we showed—using a model without fiscal shocks and with quasilinear preferences and rigid prices—that the planner does not find it optimal to provide abundant liquidity to the economy when government debt is positive, as doing so would forgo the rents associated with supplying the asset. Figure 4 shows that this property also holds in the model with curved utility, sticky prices, and fiscal shocks.

[42] With these values, the monetary policy rule satisfies the Taylor principle, while fiscal policy is passive in the sense of Leeper (1991). Moreover, ϕ_τ is close to the threshold that defines the determinacy region in the model and implies that government debt displays near unit root behavior, consistent with evidence in Marcet and Scott (2009).

Finally, note that the parameter values of ρ_i, ϕ_τ, and ϕ_π have little effect on the moments of f_θ that are consistent with the targets defined below.

[43] Both estimates (for the 4-week and 10-week yields) are relevant if the empirical counterpart of b_S includes all debt with maturities up to one quarter. The 2-basis-point target therefore falls within the range of these estimates. Our results do not depend critically on this target.

Assuming a lower target value (and hence a larger variance of f_θ) would simply widen the region of the state space in which the financial friction binds, necessitating an increase in $\overline{M}_S$. The qualitative features of the optimal portfolios highlighted below would remain unchanged.

Finally, to achieve the target, we simulate, using the auxiliary model, a shock to the ratio $\hat{b}_{S,t} - \hat{Y}_t$.

[44] We deliberately present a long sample to illustrate how optimal policy behaves across a wide region of the state space.

Long bonds fluctuate over a wide range in the simulation, from low to high levels of debt. Since short bonds remain consistently below the limit $\overline{M}_S$, this suggests that the property holds throughout the simulated state space.[45]

Note that, from a theoretical standpoint, the incentive to substitute long-term for short-term debt is even weaker in an environment with sticky prices and curved utility. Long bonds provide fiscal insurance through endogenous price fluctuations. These price adjustments make the real payoff on long-term debt state-contingent, allowing it to absorb fiscal shocks. As a result, the planner has even less incentive to expand the supply of short-term debt than in the model with quasilinear utility and rigid prices.

However, the optimal policy is also far from the one that maximizes the hedging benefits from long debt. From the discussion of subsection 4.1, we know that hedging is maximized when the government issues long debt and saves short. As we have ruled out lending to the private sector, the government might want to get as close as possible to this portfolio by setting short bonds close to zero.

This does not occur. Intuitively, in setting the optimal policy, the planner internalizes the social value of supplying the asset. The optimal policy strikes a balance between the portfolio that maximizes the social value (i.e. $b_{S,t} = \overline{M}_S$) and the one that improves fiscal outcomes through rents and hedging. Remarkably, short-term debt issuance exhibits low volatility, remaining within a narrow range (0.24–0.26).

Given the limited fluctuations in short bonds, it is not surprising that optimal policy implies a wide range of long-term debt positions. Remarkably, long-term debt behaves much as in the canonical Ramsey framework, inheriting the unit-root component of taxes.

Optimal Inflation. Turning to the simulated path of inflation shown in the bottom right panel of Figure 4, we note that, consistent with our previous results, the average inflation rate is positive. The planner sets a positive target. Inflation is positive in the sample and falls below one percent only for a few periods when a sequence of large negative spending shocks leads to a rapid reduction of debt.

Financing Spending Shocks. To more closely examine the properties of the optimal debt policy, Figure 5 displays the responses of the government portfolio to a sequence of positive spending shocks. We assume a sequence of one-standard-deviation shocks occurring in each period from period $t = 1$ to $t = 5$. Thereafter, spending reverts to its mean at rate ρ_G. [46]

Focus on the left panel, which plots the response of short-term debt. As is evident, unexpected positive shocks lead to a reduction in the supply of the short-term asset. Consistent with our theoretical results (Proposition 5), short-term debt is used by the government as a shock absorber.

[45] This also remains true in the rare instances when long-term debt approaches its lower bound, $\underline{M}_L = 0$ (not visible in the figure.)

[46] The variables are expressed as percentage deviations from their values in the absence of shocks. To generate the series (with and without shocks), the model is first simulated for 100 periods without shocks so that the initial bond quantities are close to their stochastic (risky) steady state. The debt-to-GDP ratio is 70 percent of annual GDP.

Figure 4: **Model Simulation: Baseline**

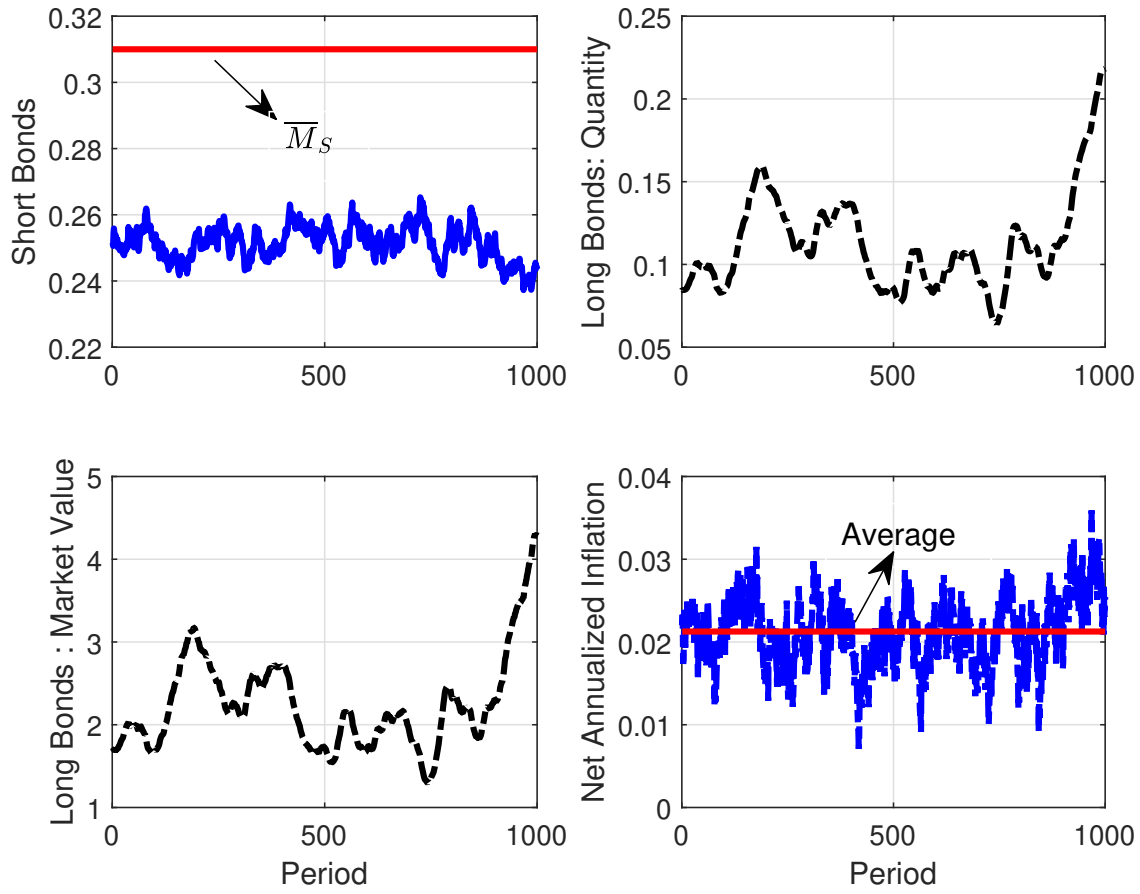

Notes: The figure shows 1000 simulated periods from the baseline optimal policy model. The top-left panel plots the quantity of short bonds. The horizontal red line is the upper bound on short-term debt. The top-right panel plots the quantity of long bonds and the bottom-left their market value. The bottom-right shows annualized net inflation rates in the model. The horizontal red line denotes the sample average.

By reducing its quantity, the government increases the rents it extracts and is able to partly finance spending through higher rents.

This combination of higher spending and lower short-term debt issuance requires a greater supply of long bonds. The figure confirms that the long bond issuance exhibits explosive dynamics, consistent with the previous finding that long bonds in the model display similar dynamics to the canonical model.[47]

In summary, the numerical results of this section convey two key messages. First, the optimal policy maintains short-term debt in the region where financial frictions bind. The model therefore precludes fluctuations between periods of abundant and scarce liquidity, validating our approach in the previous sections of deriving theoretical and numerical results based on approximations around

[47]The Appendix shows the responses of inflation and consumption to the shock sequence. Inflation and consumption evolve in a manner consistent with our theoretical framework. Specifically, inflation increases persistently in response to the shock. Private consumption drops when the shock hits and then reverts to steady state from period 5 onward. Due to the fall in consumption and the expected recovery, real long bond prices decline. This confirms that fiscal hedging is operative in the nonlinear model.

the steady state with binding financial frictions. Second, we find that the planner pursues a policy characterized by a relatively stable issuance of short-term debt, while reducing the quantity of short bonds in periods of high spending shocks and rising debt levels.

Figure 5: **Impulse responses to spending shocks**

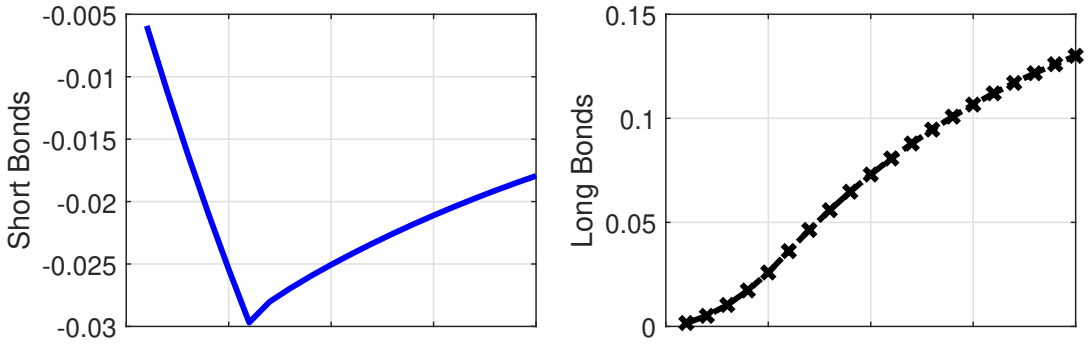

Notes: The figure shows the impulse responses of short bonds (left) and long bonds (right) to a sequence of positive spending shocks, expressed as percentage deviations from the zero-shock path.

Comparing the model with US data. We now compare the predictions of our model with U.S. data. Following the recent literature on optimal government portfolios, we use our framework to draw implications for U.S. Treasury debt management. As in this literature, our analysis in this section combines both the normative and positive aspects of the theory.

To compare the model outcomes with the data, it is necessary to specify which assets are considered liquid and safe in the U.S. economy. We follow the convention adopted in the literature, and define short-term debt as any obligation with a maturity of one year or less, rather than one quarter.[48] Long-term debt encompasses all obligations with maturities exceeding one year.

In Table 2, we report several data moments that summarize the behavior of the share of short-term debt relative to long-term debt in the U.S. economy. Specifically, we compute the sample mean of the share, the first-order autocorrelation coefficient, the standard deviation of the share, and its correlation with the debt-to-GDP ratio.

As discussed in the introduction, a large body of recent literature has sought to understand why governments issue substantial amounts of short-term debt, despite the predictions of the baseline

[48]The data used in this subsection are sourced from the OECD and cover the period starting in 1954Q3. They correspond to the dataset employed by Ghomi et al. (2026) (see that paper for further details). The dataset aggregates government debt with an original maturity of less than one year—primarily Treasury bills—which aligns with the standard definition of short-term debt used by international organizations such as the OECD, IMF, and World Bank. This measure is widely used in the literature (e.g., Faraglia et al. (2019); Debortoli et al. (2017); Krishnamurthy and Vissing-Jorgensen (2012), among others).

Recall also that though the maturity of short debt coincides with one period in the model, this assumption was mainly made for tractability. Our results could be extended to cases where short debt has longer maturity, or alternatively by interpreting the model period as one year, as in Faraglia et al. (2019); Debortoli et al. (2017).

Finally, note that Greenwood et al. (2015) employ an alternative definition of short-term bonds, which includes T-bills with maturities up to 26 weeks to isolate the "cycle component." The empirical patterns we document in this section are consistent with those reported in Greenwood et al. (2015) under this alternative definition.

debt management model, which suggests that optimal government debt should be predominantly long-term (Faraglia et al., 2019; Debortoli et al., 2017; Aparisi de Lannoy et al., 2025). These studies demonstrate that, once various frictions are introduced into the canonical framework, the optimal share of short-term debt becomes strictly positive. Our model yields a similar result.

Our theory emphasizes a different margin compared to previous contributions. In our framework, short-term bonds are valuable due to their safety and liquidity characteristics, and the government supplies these assets partly to accommodate households' demand for precautionary savings.[49] Under these assumptions, generating a positive share of short-term bonds is relatively straightforward, yet the model produces an average share that is remarkably close to its counterpart in U.S. data.

Moreover, an important implication of our model, noted in the previous paragraph, is that the issuance of short-term bonds exhibits relatively low volatility over time. The two data moments that correspond to this prediction are the standard deviation and the first-order autocorrelation coefficient of the short-term share. As shown in Table 2, the model's predictions align closely with their empirical counterparts. In particular, the standard deviation of the share is 0.024 in the data and 0.02 in the model. Moreover, both the model and the data exhibit high serial autocorrelation in the short-term debt share.

Finally, another important prediction of our model is that the issuance of short-term bonds tends to decline during periods of high spending needs or rapid increases in the debt level. Consequently, in the model, the share of short-term debt relative to long-term debt is negatively correlated with the debt-to-GDP ratio. The last row of Table 2 confirms this, with a correlation coefficient of -0.94.

In the data, the estimated correlation is -0.43. Consistent with the model, the U.S. Treasury tends to issue more long-term debt during periods when debt-to-GDP rises. However, the model implies that this correlation is somewhat too strong. Although this discrepancy can be attributed to several factors, viewed through the lens of our theory, it suggests that using short-term debt as a shock absorber—financing positive spending shocks with long-term debt only—is not fully consistent with U.S. Treasury issuance patterns.[50]

4.5 Lending Models

We now discuss how the model's implications change when the government is allowed to take negative positions in the long-term bond market.

In Appendix D, we analyze a version of the model with symmetric bounds, $\underline{M}_L = -0.3, \quad \overline{M}_L = 0.3$, which implies that the government can invest in private assets up to 167 percent of annual

[49] As noted in the introduction, this mechanism is related to the one studied by Robin Greenwood and coauthors, who highlight the role of liquid assets in Treasury issuance. Unlike their work, however, spending shocks play a central role in our analysis.

[50] In Ghomi et al. (2026) we take a close look at the US observations to identify periods in which the US government finances spending shocks using long-term debt and periods in which the shocks are financed predominately with the issuance of short bonds. We find that a large fraction of positive spending shocks are financed short-term, resulting in positive correlation between the share and the identified spending shocks. According to our model short-term financing is not optimal.

Table 2: Data and Model Outcomes

	Data	Model
Mean share	0.124	0.099
Auto-correlation	0.890	0.990
Standard deviation	0.024	0.020
Correlation with debt/GDP	-0.430	-0.940

Notes: The first column reports the mean share of short over long, the first-order serial autocorrelation coefficient, and the standard deviation of the share, as well as the correlation of the share with the debt-to-GDP ratio in the data. The second column reports the analogous moments in the optimal policy model under "No Lending". The data source is the OECD. See Ghomi et al. (2026) for variable definitions. The model statistics are computed as averages of 250 quarterly observations.

GDP. Under this specification, the optimal policy features negative long-term debt. By front-loading taxes, the government accumulates assets over the medium run.

Such asset accumulation is a standard feature of Ramsey-optimal policies (e.g., Aiyagari et al., 2002; Faraglia et al., 2019, 2016). By targeting negative debt, the government can use the interest income from its asset holdings to finance spending shocks.

Our main finding is that, in this case, the optimal policy equilibrium may feature oscillations between periods of abundant liquidity and periods in which liquidity is scarce. This result mirrors our earlier analysis showing that, when debt is negative, the optimal policy eliminates the financial friction.

However, unlike in the perfect-foresight models studied in the previous sections, the presence of shocks prevents the allocation from converging to the canonical Ramsey outcome. Instead, during periods of high spending, the government continues to rely on short-term debt as a shock absorber, pushing the economy back into a region where the financial friction binds.

In other cases, when the lower debt limit is tighter—so that the market value of long-term debt cannot fall as low as 167 percent of GDP—the qualitative findings described in the previous paragraphs continue to hold. We consider calibrations in which the government cannot lend to the private sector, or can lend only moderate amounts relative to GDP, to be the most plausible configurations of the lower-bound constraint.[51]

[51] As discussed previously, we introduced the lower-bound constraint in the model as a simple device to prevent the government from accumulating unrealistically large amounts of private assets, without explicitly modeling deeper microfoundations. For example, in Lustig et al. (2008) the motivation behind the "No Lending" assumption is that governments are unwilling to bear the idiosyncratic default risk inherent in private asset investments. During quantitative easing, the Fed did purchase private assets, and in our model the consolidated budget is the relevant equilibrium object. We therefore consider a calibration in which $\underline{M}_L$ is approximately equal to the total amount of mortgage-backed securities (MBS) purchased during QE relative to GDP. Common estimates in the finance literature indicate that these purchases corresponded to roughly 30 percent of the MBS market (Greenwood et al., 2015) by the end of QE3. During the pandemic, the Fed's MBS holdings were around 2.6–2.7 trillion, again close to 30 percent of the market. Given these figures, a plausible calibration of $\underline{M}_L$ is between -0.02 and -0.03, corresponding to approximately 12 percent of GDP. With this calibration, the model outcomes are very close to our baseline (see Appendix D).

5 Conclusion

This paper studies optimal policy in an economy in which a subset of government liabilities provides households with liquidity and safety services. Departing from the canonical Ramsey benchmark, we establish three main results. First, optimal monetary policy is characterized by a positive and time-varying inflation target. Second, the government optimally restricts the supply of liquid and safe bonds, rather than eliminating financial frictions altogether. Third, in the presence of financial frictions, optimal policy introduces volatility into the tax schedule operating through novel intertemporal distortion channels. Characterizing optimal policy in this environment is particularly challenging, because of the inherent multiplicity of optima. We develop a numerical algorithm that resolves this challenge and enables the computation of optimal policy allocations.

Our findings carry implications for the institutional design of monetary and fiscal policy. First, environments where public debt serves as a liquid/safe asset rationalize a positive time-varying inflation target. Thus, policy authorities operating in such environments may optimally tolerate higher average inflation as part of an integrated fiscal-monetary strategy. Second, debt management authorities should tilt portfolios away from liquid/safe (short-term) assets in periods of high spending and rising debt levels to increase rents from liquidity/safety provision and expand fiscal space.

Future work could revisit the questions we addressed in this paper through the lens of models with richer heterogeneity. This would allow us to incorporate heterogeneity into optimal policy design and to analyze environments in which the liquidity and safety properties of government debt arise endogenously from policy choices.

References

Acikgöz, Ö., M. Hagedorn, H. Holter, and Y. Wang (2018). The optimum quantity of capital and debt. *CEPR Discussion Paper 12952*, 1–64.

Aiyagari, S. R., A. Marcet, T. J. Sargent, and J. Seppälä (2002). Optimal taxation without state-contingent debt. *Journal of Political Economy 110*(6), 1220–1254.

Angeletos, G.-M. (2002). Fiscal policy with noncontingent debt and the optimal maturity structure. *The Quarterly Journal of Economics 117*(3), 1105–1131.

Angeletos, G.-M., F. Collard, and H. Dellas (2022). Public debt as private liquidity: Optimal policy. mimeo.

Aparisi de Lannoy, L., A. Bhandari, D. Evans, M. Golosov, and T. Sargent (2025). Managing public portfolios. *Journal of Political Economy 133*(12), 3903–3951.

Auclert, A., M. Cai, M. Rognlie, and L. Straub (2024). Optimal long-run fiscal policy with heterogeneous agents. Working paper.

Bansal, R. and W. J. Coleman (1996). A monetary explanation of the equity premium, term premium, and risk-free rate puzzles. *Journal of Political Economy 104*(6), 1135–1171.

Barro, R. J. (1979). On the determination of the public debt. *Journal of Political Economy 87*(5, Part 1), 940–971.

Bhandari, A., D. Evans, M. Golosov, T. Sargent, et al. (2019). The optimal maturity of government debt. mimeo.

Bigio, S., G. Nuño, and J. Passadore (2023). Debt-maturity management with liquidity costs. *Journal of Political Economy Macroeconomics 1*(1), 119–190.

Buera, F. and J. P. Nicolini (2004). Optimal maturity of government debt without state contingent bonds. *Journal of Monetary Economics 51*(3), 531–554.

Canzoneri, M., R. E. Cumby, and B. T. Diba (2016, October). Optimal money and debt management: Liquidity provision vs. tax smoothing. *Journal of Monetary Economics 83*, 39–53.

Chafwehé, B., C. de Beauffort, and R. Oikonomou (2026). Managing the inflation-output trade-off with public debt portfolios. *Journal of Economic Dynamics and Control 183*(February), 473–501.

Challe, E., F. Le Grand, and X. Ragot (2013). Incomplete markets, liquidation risk, and the term structure of interest rates. *Journal of Economic Theory 148*(6), 2483–2519.

Dávila, E. and A. Schaab (2023). Optimal monetary policy with heterogeneous agents: Discretion, commitment, and timeless policy. *NBER Working Paper 30961*, 1–47.

Debortoli, D., R. Nunes, and P. Yared (2017). Optimal time-consistent government debt maturity. *The Quarterly Journal of Economics 132*(1), 55–102.

Den Haan, W. J. and A. Marcet (1990). Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics 8*(1), 31–34.

Faraglia, E., A. Marcet, R. Oikonomou, and A. Scott (2013). The impact of debt levels and debt maturity on inflation. *The Economic Journal 123*(566), F164–F192.

Faraglia, E., A. Marcet, R. Oikonomou, and A. Scott (2016). Long term government bonds. mimeo.

Faraglia, E., A. Marcet, R. Oikonomou, and A. Scott (2019). Government debt management: The long and the short of it. *The Review of Economic Studies 86*(6), 2554–2604.

Faraglia, E., A. Marcet, and A. Scott (2010). In search of a theory of debt management. *Journal of Monetary Economics 57*(7), 821–836.

Geromichalos, A., L. Herrenbrueck, and S. Lee (2023). Asset safety versus asset liquidity. *Journal of Political Economy 131*(5), 1172–1212.

Geromichalos, A., J. Lee, and K. Oikawa (2016). Over-the-counter trade and the value of assets as collateral. *Economic Theory 62*(3), 443–475.

Ghomi, M., J. Mankart, R. Priftis, and R. Oikonomou (2026). Debt maturity and governme6t spending multipliers. mimeo.

Gorton, G. and A. Metrick (2012). Securitized banking and the run on repo. *Journal of Financial Economics 104*(3), 425–451.

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

Hagedorn, M. (2018). A demand theory of the price level. mimeo.

Hagedorn, M. (2024). The failed theory of the price level. *CEPR Discussion Paper 18786*, 1–40.

Judd, K. L., L. Maliar, and S. Maliar (2011). Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models. *Quantitative Economics 2*(2), 173–210.

Judd, K. L., L. Maliar, and S. Maliar (2015). Merging simulation and projection approaches to solve high-dimensional problems. *Quantitative Economics 6*(1), 1–47.

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

Le Grand, F., A. Martin-Baillon, and X. Ragot (2026). Should monetary policy care about redistribution? optimal monetary and fiscal policy with heterogeneous agents. *The Review of Economic Studies (forthcoming)*, rdag019.

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

Leeper, E. M. and X. Zhou (2021). Inflation's role in optimal monetary-fiscal policy. *Journal of Monetary Economics 124*, 1–18.

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

Lustig, H., C. Sleet, and Ş. Yeltekin (2008). Fiscal hedging with nominal assets. *Journal of Monetary Economics 55*(4), 710–727.

Maliar, L. and S. Maliar (2003). Parameterized expectations algorithm and the moving bounds. *Journal of Business & Economic Statistics 21*(1), 88–92.

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

Marcet, A. and A. Scott (2009). Debt and deficit fluctuations and the structure of bond markets. *Journal of Economic Theory 144*(2), 473–501.

Mian, A., L. Straub, and A. Sufi (2025). A goldilocks theory of fiscal deficits. *American Economic Review 115*(12), 4253–4291.

Nosbusch, Y. (2008). Interest costs and the optimal maturity structure of government debt. *The Economic Journal 118*(527), 477–498.

Piazzesi, M., M. Schneider, P. Benigno, and J. Y. Campbell (2007). Equilibrium yield curves [with comments and discussion]. *NBER macroeconomics Annual 21*, 389–472.

Rotemberg, J. J. (1982). Sticky prices in the united states. *Journal of Political Economy 90*(6), 1187–1211.

Rudebusch, G. D. and E. T. Swanson (2012). The bond premium in a DSGE model with long-run real and nominal risks. *American Economic Journal: Macroeconomics 4*(1), 105–43.

Schmitt-Grohé, S. and M. Uribe (2004a). Optimal fiscal and monetary policy under imperfect competition. *Journal of Macroeconomics 26*(2), 183–209.

Schmitt-Grohé, S. and M. Uribe (2004b). Optimal fiscal and monetary policy under sticky prices. *Journal of Economic Theory 114*(2), 198–230.

Schneider, A. (2022). Risk-sharing and the term structure of interest rates. *The Journal of Finance 77*(4), 2331–2374.

Sims, C. A. (2026). Optimal fiscal and monetary policy with distorting taxes. *American Economic Journal: Macroeconomics 18*(1), 1–23.

Siu, H. E. (2004). Optimal fiscal and monetary policy with sticky prices. *Journal of Monetary Economics 51*(3), 575–607.

Williamson, S. D. (2016). Scarce collateral, the term premium, and quantitative easing. *Journal of Economic Theory 164*, 136–165.