Optimal Taxation with Incomplete Markets–an Exploration via Reinforcement Learning *

Zhigang Feng[†] Jiequn Han[‡] Shenghao Zhu[§]

December 11, 2023

Abstract

In this paper, we delve into the optimal taxation, government debt/asset management, and price volatility in economies with incomplete markets. Employing a quasi-linear preference framework, we find that any finite level of government asset is sustainable, and the derivative of the value function behaves as a martingale, exhibiting no convergence. However, when taxes or assets are bounded, the asset accumulation process assumes an ergodic nature, even as government expenditure adheres to a Markov process. Moreover, we demonstrate that the well-known Martingale convergence result in existing literature is attributed to a weak precautionary savings motive by the government. Conversely, with a sufficiently strong precautionary savings motive, the asset accumulation process regains its ergodic characteristics. Additionally, we introduce a computational method based on machine learning to solve the equilibrium set of endogenous state variables, policy, and value functions. Our numerical simulations reveal that optimal taxation and government debt/asset management are closely intertwined, exerting substantial influence on asset pricing.

JEL Classification: F3, F4.

Keywords: Ramsey Taxation, Asset Pricing, Debt Management, Deep Neural Network, Reinforcement Learning.

---

*We thank Felix Kubler, Tom Sargent, John Stachurski, Yi Wen for their insightful comments. We also extend our thanks to the participants of seminars across various departments for their valuable contributions to our discussions. This work was completed utilizing the Holland Computing Center of the University of Nebraska, which receives support from the Nebraska Research Initiative, as well as Anvil at Purdue through allocation SOC230021 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

[†]Department of Economics, University of Nebraska, Omaha, NE 68106. E-mail: z.feng2@gmail.com

[‡]Center for Computational Mathematics, Flatiron Institute, New York, NY 10010. Email: jiequnhan@gmail.com.

[§]Department of Economics, University of International Business and Economics, Bejing, China. E-mail: zhushenghao@yahoo.com

1 Introduction

In this paper, we delve into the intricate relationship between optimal taxation and the management of government debt and assets in economies characterized by incomplete markets and full governmental commitment. We focus on a closed economy subject to exogenous public spending shocks. In such an environment, the government predominantly relies on labor income taxes and a one-period non-contingent bond, purchased by domestic households, to finance unexpected expenditures.

This setting naturally leads the government to preemptively manage tax distortions through asset accumulation. These assets are strategically utilized to finance future public spending, either partially or entirely, through interest income. At the household level, debt accumulation is influenced by expected future disposable incomes, which are, in turn, shaped by government policy measures.

Consequently, governmental decisions about taxation, debt issuance, and asset management are deeply interconnected, exerting a significant impact on one another. This interplay crucially affects asset prices. Fiscal policy's influence on household borrowing and saving behavior becomes a key determinant of both the marginal rate of substitution and bond prices. These complex interactions illustrate how bond prices and interest rates can affect the government's capacity to manage tax distortions. Our insights highlight the nuanced relationship between government policy, household finance, and economic stability, offering a comprehensive view of fiscal strategy development and implementation within a broader economic framework.

Our analysis starts with a model based on quasi-linear preferences, where the marginal utility of consumption is constant, setting the bond price as a fixed entity. We initially consider a scenario with no constraints on the government's ability to adjust the tax rate, borrow, or save.

When the government has access to lump-sum transfers, as Aiyagari et al. (2002) suggests, it can easily achieve a first-best scenario. The government would accumulate assets equal to the present value of the highest foreseeable public spending. This asset base allows the government to finance future expenditures through interest income, maintaining solvency even during the worst G shocks. This approach is akin to the precautionary saving motive in household income fluctuation models. In favorable spending shocks, the government redistributes surplus back to households, thus balancing the budget.

Without lump-sum transfers, the government may struggle to maintain surplus interest earnings from assets during favorable G shocks. This limitation impedes its asset accumulation capacity, affecting its ability to finance public expenses through interest income during adverse expenditure shocks. A remedy to this is to lend the surplus to households, limited by their debt capacity, which hinges on future net-tax earnings and prevailing tax policies.

To explore the interplay between tax policy and asset/debt accumulation, we examine three scenarios: i) no limits on government taxation or asset/debt; ii) unrestricted taxation with a defined asset limit; iii) a predefined tax rate limit with no asset/debt constraints. Our findings show that the behavior of the value function's derivative varies with the constraints on government assets and debts. In the absence of restrictions, the derivative behaves as a martingale with no convergence. However, with a lower bound on government assets, it becomes a super-martingale, potentially diverging to negative infinity.

In the first scenario, the government achieves a first-best allocation through a Ponzi scheme, allowing it to balance its budget with a zero tax rate. This scheme is sustained by the unbounded tax policy, which induces extended labor hours, ensuring any required tax revenue for future Ponzi scheme closure.

With no tax rate restriction, any arbitrary asset limit impedes the government's ability to absorb surplus interest income by lending to households. Regardless of the asset limit level, the economy occasionally hits this threshold, typically after successive favorable public expenditure shocks. The asset constraint also sets a strict minimum for tax rates. Additionally, we found that the equilibrium bond process has a unique invariant distribution, significantly extending Proposition 1 of Bhandari et al. (2016).

Once we incorporate a tax constraint into our model, the policy inherently shapes household income, as dictated by the household's intra-temporal optimality condition. Once the post-tax earnings are capped, the household's borrowing potential gets similarly restricted. As we consider a closed economy setting, households and government are the only trading partners in bond market. This dynamic naturally sets threshold for government bond holdings, defining the so-called natural asset and debt limits.

In equilibrium, both the government's asset limit and the household's debt limit are endogenously determined, unless there are specific constraints on assets or debt. Moreover, when utility functions account for income effects, bond prices become endogenous as well. The interest rate might show a correlation with the tax rate (τ), which complicates the calculation of the household's natural debt limit. Furthermore, the state space (Ω) includes an additional endogenous variable: the promised marginal utility (λ) for the household. These factors must all be characterized in tandem with the equilibrium policy and value functions, which presents a significant challenge. This is because these functions are well-defined only in relation to a properly specified set Ω.

To overcome these challenges, we have devised an innovative numerical algorithm that builds upon the self-generating method presented in Feng et al. (2014) and Feng (2015), specifically tailored for delineating the endogenous state space. Additionally, we incorporate recent developments in resolving dynamic equilibrium models through machine learning, as demonstrated in Han et al. (2022). This approach involves employing deep neural networks for precise and efficient approximation of high-dimensional non-linear functions.

Our novel approach is distinguished by our utilization of a deep neural network to represent the unknown set of equilibrium state variables. We subsequently reformulate the government's problem as an unconstrained optimization problem, introducing two key modifications. First, we supplement the objective function with an extra term that encapsulates the admissibility constraint concerning state variables. This constraint necessitates that both present and future state variables belong to the same candidate equilibrium set. Secondly, we substitute the agent's choice variables with the government's policy instrument by applying the agent's first-order condition.[1] Additionally, our method leverages recent advancements

[1]Yang and Zhu (2022) employed this method to a Ramsey taxation problem within a heterogeneous-

in reinforcement learning by solving this optimization problem through a two-step process: sequential updating of the policy and value functions for a given set of equilibrium state variables. This methodology is not only straightforward to implement but also scales efficiently for problems of higher dimensions.

From a machine learning perspective, especially within the context of nonparametric learning, the iterative process of searching for the policy function can be perceived as nonparametric estimation. Chen and White (1998) utilized nonparametric adaptive learning with feedback for estimating the policy function in the context of dynamic programming. The introduction of a feedback rule to the stochastic approximation method addresses the forward-looking characteristic of dynamic economic problems. However, the dynamic Stackelberg game, as in our model, introduces an additional challenge, which is determining the domain of the policy functions. To circumvent this issue, we use a deep neural network (DNN) to approximate the equilibrium set of endogenous state variables, which transforms the domain-finding issue into a binary classification problem (0 for out of the domain and 1 for in the domain).[2] Thus, our objective is to estimate these functions non-parametrically as the solution to a fixed-point problem.

We solve the Ramsey equilibrium and conduct quantitative analyses to understand the effects of tax and asset limits on economic dynamics. Numerical results show that when there is a sufficiently high lower bound on tax rate, or on the government's asset limit, the economy will rebound upon reaching this boundary. However, these lower bounds (whether tax rate or asset limit) become an absorbing state when the values for these bounds are set sufficiently low.

Our numerical analyses also reveal the existence of a dynamic Laffer curve for debt accumulation and a hump-shaped relationship between the government's debt limit and the promised marginal utility, λ. When starting with a sufficiently large asset position, the government might manage to command near-zero, or even negative, interest rates to alleviate the household's debt obligation. This offers justification for the sustainability of initial asset/debt positions. As a consequence, the asset price may exhibit volatility that is significantly larger than what would be suggested by the primitive characteristics of the economy.

Related literature This paper studies the same economic environment as Aiyagari et al. (2002). While they address complexity in the characterization of the endogenous asset/debt limits via lump-sum transfer, we characterize such limits based on an innovate computational algorithm. Compared with Bhandari et al. (2017), which use approximation method around the steady state, our paper finds the global solution of the Ramsey taxation problem. We find that there may exist multiple steady states, which suggest that the perturbation around the steady state may be problematic. Another advantage of our approach is that it can generate large volatility in asset prices as observed in the data. The link with Cho et al.

agents model.

[2] Renner and Scheidegger (2020) develop a similar method for the computation of discrete-time dynamic incentive problems.

(2002).

Jiang et al. (2022) study the Ramsey taxation problem using a continuous-time model. Both Jiang et al. (2022) and our paper solve the full-commitment solution. Our paper has a discrete-time model, and the promise variable of the next period is the choice variable in the government in the current period. Thus, we can solve the promise choice in the dynamic programming. In the continuous-time model, one has to introduce the choice variable corresponding to the promise variable into the Ramsey problem.

The paper is organized as follows. Section 2 sets out the basic model. Section 3 derives some analytical properties for the Ramsey government's problem. Section 4 explains the parameterization for numerical examples and the computational methods for solving the equilibrium. In section 5 we study the equilibrium via numerical simulation. Section 6 concludes. In Appendix we outline proofs of our main results.

2 The Economy

Time is discrete and indexed by $t = 0, 1, ...$ The economy is populated by a measure one of continuum of households. The representative household starts the economy with b_0 units of bond issued by the government and supply labor, l_t, at every date $t = 0, 1, \cdots$.

The aggregate uncertainty is represented by z_t, which affects the level of public spending G_t. There is a linear production technology as in Lucas and Stokey (1983). The output can be used for private and public consumption c_t, G_t. The government taxes income at a flat rate, τ_t, to finance public consumption, G_t, and can borrow a one-period bond from the households at price q_t.

2.1 The household

For a given sequence of taxes and public consumption $\{\tau_t, G_t\}_{t=0}^{\infty}$ and bond prices $\{q_t\}_{t=0}^{\infty}$, the representative household chooses an optimal plan for private consumption, hours worked, and bond holding $\{c_t, l_t, b_{t+1}\}_{t=0}^{\infty}$ to maximize the inter-temporal utility function,

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u\left(c_t\right) - h(l_t)\right] \tag{1}$$

subject to the income and the resource constraint,

$$(1 - \tau_t)l_t + b_t \geq q_t b_{t+1} + c_t, \tag{2}$$

for $t \geq 0$, and the no-Ponzi-game condition,

$$\lim_{t\to\infty} \mathbb{E}_0 \left[\left(\prod_{i=1}^{t} q_i\right) b_{t+1}\right] \leq 0, \tag{3}$$

for initial b_0. $h(l_t)$ represents the dis-utility of labor.

One-period utility functions u and h are assumed to satisfy standard regularity conditions over $c_t \geq 0$, and $l_t \geq 0$. Without any further constraints, the optimal behavior of the household can be characterized by the first-order conditions,

$$u'(c_t)q_t = \beta \mathbb{E}_t u'(c_{t+1}), \tag{4}$$

and

$$u'(c_t)(1-\tau_t) = h'(l_t), \tag{5}$$

for $t \geq 0$.

2.2 The government

The government's problem is to choose a contingent plan for the tax rate, and the quantity of bond holdings $\{\tau_t, B_{t+1}\}_{t=0}^{\infty}$ to maximize the social welfare,

$$W(z_0, B_0) = \max_{\{\tau_t, B_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u\left(c_t\right) - h(l_t)\right] \tag{6}$$

subject to the period-by-period budget constraint,

$$G(z_t) + B_t \leq \tau_t l_t + q_t B_{t+1}, \tag{7}$$

for $t \geq 0$, and the no-Ponzi-game condition,

$$\lim_{t \to \infty} \mathbb{E}_0 \left[\left(\prod_{i=1}^{t} q_i \right) B_{t+1} \right] \leq 0, \tag{8}$$

over all equilibrium paths $\{c_t, l_t, b_{t+1}\}_{t=0}^{\infty}$, and for initial (z_0, B_0). Note that each equilibrium path represents an optimal response from households to the policy choices made by the government.

2.3 Recursive representation of the government's problem

Definition 2.1. *For given (B_0, z_0), a competitive equilibrium is a set of government policies $\Lambda = \{\tau_t, B_{t+1}\}_{t=0}^{\infty}$, bond prices $\{q_t\}_{t=0}^{\infty}$, and consumer choices $\{c_t, l_t, b_{t+1}\}_{t=0}^{\infty}$ such that*

1. *Λ satisfies budget constraint (7) and (8) at all times;*

2. *$\{c_t, l_t, b_{t+1}\}_{t=0}^{\infty}$ solves the optimization problem (1) subject to (2) and (3), for the given policy Λ and bond prices $\{q_t\}_{t=0}^{\infty}$;*

3. $b_t + B_t = 0$ *and* q_t *satisfies (4) for all* $t \geq 0$.

Definition 2.2. *For given* (z_0, B_0), *the Ramsey government chooses the optimal* $\Lambda = \{\tau_t, B_{t+1}\}_{t=0}^{\infty}$ *to maximize the social welfare as defined in (6), over the set of competitive equilibria.*

As in Kydland and Prescott (1980), we recast the government's problem recursively by including the promised shadow value of investing in government bond as a state variable. Specifically, we define $\lambda = u'(c)$ and write the Ramsey government's problem at $t > 0$ recursively as follows.

(P-1)

$$V(z, B, \lambda) = \max_{(\lambda_+(z_+))} \{u(c) - h(l) + \beta \cdot \mathbb{E}V(z_+, B_+, \lambda_+)\} \tag{9}$$

subject to

$$G(z) + B \leq \tau l + qB_+ \tag{10}$$

$$c + G(z) = l \tag{11}$$

$$u'(c) = \lambda \tag{12}$$

$$(1 - \tau)u'(c) = h'(l) \tag{13}$$

$$qu'(c) = \beta\mathbb{E}[\lambda_+|z] \tag{14}$$

$$(B, \lambda) \in \Omega(z), (B_+, \lambda_+) \in \Omega(z_+), \tag{15}$$

where $\Omega(z)$ denotes the endogenous equilibrium feasible set for B and λ, which is unknown and will be characterized via an iterative procedure detailed in Section 4.[3] Note that equation (10) is the government's budget constraint, equation (11) is the resource constraint for the economy, equation (12) is the promise keeping constraint from date-0 Ramsey government. Equation (13) gives the tax rate and equation (14) yields the price for bond. Constraint (15) requires that bond issuance and the promised shadow prices for bond must lie in the endogenous feasible set.

The date-0 government sets the policy for all future governments and isn't subject to the promise keeping constraint (12), and has a slightly different problem to solve,

(P-2)

$$W(z_0, B_0) = \max_{(\lambda_+(z_+))} \{u(c) - h(l) + \beta \cdot \mathbb{E}V(z_+, B_+, \lambda_+)\} \tag{16}$$

[3]Given the state variable values, we determine c by utilizing equation (12), l through equation (11), and τ with the aid of equation (13). Furthermore, our choices of $\{\lambda_+(z_+)\}$ in conjunction with equation (14) yield q, which subsequently helps us solve for B_+ using equation (10).

subject to

$$
\begin{aligned}
G(z) + B &\le \tau l + qB_+ \\
c + G(z) &= l \\
(1-\tau)u'(c) &= h'(l) \\
qu'(c) &= \beta \mathbb{E}[\lambda_+ | z] \\
(B_+, \lambda_+) &\in \Omega(z_+).
\end{aligned}
$$

3 Analytical Results

In this section, we conduct a thorough examination of the fundamental theoretical aspects of the baseline economy. Our initial focus is on delineating the endogenous constraints that impact government assets and debt. Following this, we delve into an analysis of the long-term dynamics of aggregate variables and the implications for government policies.

3.1 Endogenous asset and debt limits

Each period the government is subject to a budget constraint, denoted as (7). By multiplying the period $t+1$ constraint by q_t, then summing the period t and $t+1$ constraints, we deduce the following inequality:

$$
B_t \le \tau_t \cdot l_t - G_t(z_t) + q_t \cdot \tau_{t+1} \cdot l_{t+1} - q_t \cdot G_{t+1}(z_{t+1}) + q_t \cdot q_{t+1} \cdot B_{t+2}. \tag{17}
$$

Following this, we multiply the period $t+2$ constraint by $q_t \cdot q_{t+1}$ and combine it with equation (17). Repetition of this operation for $t \to \infty$ yields the subsequent inequality:

$$
B_t \le \tau_t \cdot l_t - G_t(z_t) + \sum_{i=t+1}^{\infty} \left\{ \Pi_{j=t}^{i} q_j \cdot (\tau_{j+1} \cdot l_{j+1} - G_{j+1}(z_{j+1})) \right\} + \lim_{T\to\infty} \Pi_{j=t}^{T} q_j \cdot B_{T+1}. \tag{18}
$$

In a context where preferences are quasi-linear, the bond price remains constant ($q_t = \beta$). Let define $\bar{R}$ as the peak total tax revenue ($\tau_t l_t$) on the Laffer curve. The government's debt limit is then given by $\bar{B} = \frac{\bar{R} - \max G}{1-\beta}$. In scenarios where the government can utilize lump-sum transfers, the asset limit is established as $\underline{B} = -\frac{\max G}{1-\beta}$. As posited by Aiyagari et al. (2002), the government is expected to cover all its expenditures through the interest generated from its assets, as expressed by $q_t \cdot \underline{B} - \underline{B} = \max G$; thus, the government does not need additional taxation, thereby minimizing economic distortions. However, this asset limit tightens when the government's capacity to implement lump-sum transfers is impaired.

Consider a scenario where the government experiences a more favorable G shock, effectively leading to an excess in revenue from interest payments to the accrued assets. Specifically, the government might find itself with a surplus of $(\max G - G_t) > 0$, which exceeds the government's budget constraint. Under these circumstances, the government has the option to return this surplus to the household through a lump-sum transfer. When lump-sum transfers are not an option, the government has two alternatives for managing a budget surplus: lending it to households or implementing a negative tax. The feasibility of these strategies is contingent on the flexibility of the tax rate, which also influences the determination of the asset limit. For instance, if the tax rate cannot drop below zero, the government is unable to redistribute any surplus resulting from the G shock. To avoid this predicament, the government might choose to limit the accumulation of assets. Furthermore, the minimum limit of the tax rate is crucial as it dictates the maximum net-tax income available to households, impacting their capacity to borrow. In a closed economy, it's important to note that the government's assets are essentially the household's debt.

When considering more general preferences, there are several challenges associated with determining the asset and debt limits. Firstly, bond prices, tax rates chosen by the Ramsey government, and the private sector's responses to taxation are all endogenous and must be determined in equilibrium. Secondly, the households' optimization problem needs to be well-defined over the specified asset and debt limits. Additionally, the domain for the endogenous state variable λ in (P-1) remains unknown and fluctuates with the value of B and the shock z.

In order to address these challenges, Section 4 introduces a computational method, grounded in machine learning, that identifies the set Ω through an iterative process. This set determines the endogenous asset and debt limits along with the domain for the promised shadow value of the bond λ. For the remainder of this section, we assume that the set Ω has been established to simplify the exposition.

3.2 Characterization of equilibrium

3.2.1 The quasi-linear case

Our analysis begins with a quasi-linear preference represented by $u(c) - h(l) = c - h(l)$. Let us consider the function

$$h(l) = \gamma_l \frac{l^{1+\chi}}{1+\chi}$$

where $\chi > 0$. Given this framework, the expected shadow value of the bond is a constant and can be represented as $\lambda = 1$. We also assume that the set of possible values for G_t is denoted as $\mathcal{G}$, and is defined as

$$\mathcal{G} \equiv \{G^1, G^2, \ldots, G^n\},$$

where

$$0 < G^1 < G^2 < \ldots < G^n.$$

For any G within $\mathcal{G}$, the sum of transition probabilities

$$\sum_{G_+} \pi(G_+|G) = 1.$$

Additionally, $\pi(G_+|G) > 0$ for every pair (G, G_+) within the Cartesian product $\mathcal{G} \times \mathcal{G}$.

With no constraints on B or τ To streamline the analysis, we define the total tax revenue for the government as $Z = \tau l$, where τ represents the tax rate and l denotes labor income. From the equation (13), we obtain $l = \left(\frac{1-\tau}{\gamma_l}\right)^{\frac{1}{\chi}}$. This gives us the relationship $Z(\tau) = (\gamma_l)^{-\frac{1}{\chi}} \tau (1-\tau)^{\frac{1}{\chi}}$. Differentiating $Z(\tau)$ with respect to τ, we get

$$Z'(\tau) = (\gamma_l)^{-\frac{1}{\chi}} \left(1 - \frac{1+\chi}{\chi}\tau\right) (1-\tau)^{\frac{1}{\chi}-1}. \tag{19}$$

From this, it is evident that there exists a critical value $\bar{\tau} = \frac{\chi}{1+\chi}$ which is less than 1. For values of τ less than $\bar{\tau}$, $Z'(\tau) > 0$ and for $\bar{\tau} < \tau < 1$, $Z'(\tau) < 0$. Further, let's define $\bar{Z} = (\gamma_l)^{-\frac{1}{\chi}} \bar{\tau}(1-\bar{\tau})^{\frac{1}{\chi}}$, which is greater than 0. For all τ less than or equal to 1, we have $Z \leq \bar{Z}$. With these derivations, the government's optimization problem can be reformulated as follows:

$$\max_{\{\tau_t, B_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\left(\frac{1-\tau_t}{\gamma_l}\right)^{\frac{1}{\chi}} \frac{\tau_t + \chi}{1+\chi} - G_t \right] \tag{20}$$

subject to

$$G_t + B_t \leq (\gamma_l)^{-\frac{1}{\chi}} \tau_t (1-\tau_t)^{\frac{1}{\chi}} + \beta B_{t+1}. \tag{21}$$

Since $Z'(\tau) > 0$ for $\tau < \bar{\tau}$, we can define the function $\tau(Z)$ such that $Z = (\gamma_l)^{-\frac{1}{\chi}} \tau(Z) \left[1 - \tau(Z)\right]^{\frac{1}{\chi}}$ for $Z \leq \bar{Z}$. Now, let us define another function: $\Gamma(Z) = \left(\frac{1-\tau(Z)}{\gamma_l}\right)^{\frac{1}{\chi}} \frac{\tau(Z)+\chi}{1+\chi}$. Consequently, the derivatives are given by:

$$\Gamma'(Z) = \frac{\tau}{(1+\chi)\tau - \chi}, \tag{22}$$

and

$$\Gamma''(Z) = -\frac{\chi \tau'(Z)}{[(1+\chi)\tau - \chi]^2}, \tag{23}$$

for $Z < \bar{Z}$. Given that $\tau'(Z) = \frac{(\gamma_l)^{\frac{1}{\chi}}}{(1-\tau)^{\frac{1}{\chi}-1}\left(1 - \frac{1+\chi}{\chi}\tau\right)} > 0$ for $Z < \bar{Z}$, it implies $\Gamma''(Z) < 0$ for $Z < \bar{Z}$. This means that $\Gamma(Z)$ is strictly concave. The government's optimization problem can be written as:

$$\max_{\{Z_t, B_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\Gamma(Z_t) - G_t\right] \tag{24}$$

subject to

$$G_t + B_t \leq Z_t + \beta B_{t+1}. \tag{25}$$

The corresponding Euler equation is

$$\Gamma'(Z_t) = \mathbb{E}_t \Gamma'(Z_{t+1}). \tag{26}$$

Based on equation (22), it's evident that $\Gamma'(Z) \leq \frac{1}{1+\chi}$. Let's define $\Psi_t = \frac{1}{1+\chi} - \Gamma'(Z_t) \geq 0$. Then, we observe that:

$$\Psi_t = \mathbb{E}_t \Psi_{t+1}. \tag{27}$$

Based on the relationship $\Psi_0 = \frac{1}{1+\chi} - \Gamma'(Z_0) = \frac{1}{1+\chi} - \frac{\tau_0}{(1+\chi)\tau_0 - \chi}$, we deduce Ψ_0 is finite. Thus, the process $\{\Psi_t\}_{t=0}^{\infty}$ behaves as a martingale. By invoking the Martingale Convergence Theorem, we can infer the existence of a random variable Ψ_∞ such that $\mathbb{E}(\Psi_\infty) = \Psi_0$ and $\lim_{t\to\infty} \Psi_t = \Psi_\infty$ almost surely (a.s.). Analogously, there must exist a random variable Z_∞ for which

$$\lim_{t\to\infty} Z_t = Z_\infty \quad a.s. \tag{28}$$

However, it's worth noting that Z_∞ could assume the value $-\infty$. It's important to mention that there are no constraints on the value of τ, except that the government might prefer not to choose τ greater than $\bar{\tau}$. Consequently, the process $\{Z_t\}_{t=0}^{\infty}$ possesses an infinite number of stationary distributions. Each of these distributions is contingent on the initial condition Z_0. Furthermore, we have the following proposition characterizes the behavior of B_t in the long run.

Proposition 3.1. $\Pr(|B_t| > M$ *infinitely often*$) > 0$ *for any* $M > 0$.

The Ramsey government strategically aspires to accumulate a sufficient level of assets. This accumulation is aimed at ensuring that the interest income is adequate to cover exogenous public expenditures, even under the least favorable conditions. If a favorable expenditure shock occurs, the government will generate a surplus from interest income. This surplus can be returned as a lump-sum transfer, as described by Aiyagari et al. (2002). By adopting this approach, the government seeks to minimize tax-induced distortions, thereby aligning closely with the principles of optimal taxation to attain a 'first-best' scenario in economic efficiency. A more intriguing scenario arises when lump-sum transfers are ruled out. In this situation, the government may lend the surplus to the households, translating into household debt. Given that G follows an $AR(1)$ process, this lending can recur, causing government assets (and correspondingly, household debt) to potentially accumulate indefinitely. If the tax rate τ is allowed to become significantly negative, this can be a sustainable equilibrium outcome. The equation (13) suggests that labor supply and household income would then reach a level adequate for repaying the government upon request. Consequently, the endogenous debt limit expands significantly. For the analysis conducted thus far, it is important to emphasize that there are no imposed constraints on either τ or B.

With an ad hoc bound on B If government bond issuance is bounded, the Bellman equation of the government's problem becomes

$$V(z, B) = \max_{Z \in (-\infty, \bar{Z}], B_+ \in [\underline{B}, \bar{B}]} \{\Gamma(Z) + \beta \mathbb{E}[V(z_+, B_+)]\} - G(z)$$

subject to

$$G(z) + B \leq Z + \beta B_+.$$

From Theorem 9.6 of Stokey et al. (1989), we know that $V(z, B)$ is continuous in B. From Theorem 9.8 of Stokey et al. (1989), we know that $V(z, B)$ is strictly concave in B.

Proposition 3.2. *$V(z, B)$ is continuously differentiable in $B \in [\underline{B}, \bar{B}]$. $V_1(z, B) = \Gamma'(Z(z, B))$ for $B \in [\underline{B}, \bar{B}]$.*

Proposition 3.3. *The Euler equation of the government's problem is*

$$V_1(z, B) \geq \mathbb{E}[V_1(z_+, B_+)|G], \quad \textit{with equality if } B_+ > \underline{B}.$$

$B_+(z, B)$ is continuous and increasing in B.

$Z(z, B)$ is continuous and increasing in B. $\tau(z, B)$ is continuous and increasing in B. $l(z, B)$ is continuous and decreasing in B. $c(B, G)$ is continuous and decreasing in B.

Proposition 3.4. *There exists $G \in \mathcal{G}$ such that $B_+(z, \underline{B}) > \underline{B}$. Lower bound $\underline{B}$ is a reflecting barrier.*

Proposition 3.5. *For any B_0, there is a positive probability such that lower bound $\underline{B}$ binds in finite periods.*

Proposition 3.6. *The process $\{B_t\}_{t=0}^{\infty}$ has a unique invariant distribution.*

Proposition 3.6 extends Proposition 1 of Bhandari et al. (2016), where they presuppose that government expenditures are independent and identically distributed (i.i.d.). Foss et al. (2018) employ a regenerative concept to demonstrate the stability of some stochastic processes in economic models where the driving sequences are not necessarily i.i.d. Utilizing these results, Foss et al. (2018) extend the ergodicity results of the Brock-Mirman model from i.i.d. shocks to a Markov chain. Similarly, we extend the ergodicity result of the Ramsey problem with i.i.d. shocks to the environment of Markov chains.

The intuition behind these results rest on the regenerative concept. Starting anywhere in the state space, the process $\{B_t\}_{t=0}^{\infty}$ has a positive probability of visiting the lower bound within a finite number of periods. The lower bound of the state space is a positive recurrent atom and plays the role of a reflecting barrier. Every time the debt process hits the lower bound, the process initiates a new cycle. These cycles are i.i.d.

This ergodic result is particularly crucial for anyone aiming to solve the Ramsey problem using the perturbation method, which assumes the existence of a stochastic steady state and approximates the equilibrium around such a steady state.

The ergodicity of the bond process is contingent on the monotonicity of the policy function, which in turn is derived from the properties of the value function $V(z, B)$. The form of the value function $V(z, B)$ relies on the constraints imposed on government policies. Different constraints yield varying results. In some cases, we observe Martingale convergence. In a scenario devoid of constraints, it's optimal for the government to engage in a Ponzi scheme, thus enabling the achievement of the first-best allocation. This classification may aid in understanding recent papers on optimal taxation, which report diverse findings based on different constraints applied to government policy.

Doob's Martingale convergence theorem, c.f. Billingsley (2012), is an instrumental tool employed to examine precautionary savings in dynamic models. We discover that the Martingale convergence result arises from the fact that the lower bound of the state space acts as an absorbing barrier. The first-best allocation is achieved once the bond process is absorbed into the lower bound. In this scenario, uncertainty is eliminated and risk is extinguished.

Additionally, we ascertain that the characteristics of the lower bound of the state space are a crucial determinant in classifying the dynamics of the bond process. The lower bound can exist in one of two exclusive states–either as a reflecting barrier or an absorbing barrier. The distinction between these states represents the intensity of the government's precautionary savings. If the motive for precautionary savings is sufficiently strong, the lower bound serves as a reflecting barrier and the bond process is ergodic. Conversely, if the motivation for precautionary savings is weak, the lower bound becomes an absorbing barrier, and the bond process exhibits Martingale convergence.

Volatility Given quasi-linear preferences, the bond price stabilizes to a constant value. Consequently, the debt limit can be represented as:

$$B_t = \sum_{i=0}^{\infty} \beta^i (\tau_{t+i} l_{t+i} - G_{t+i})$$

From this representation, we can deduce the following proposition regarding the second-order moments of the economy:

Proposition 3.7. *The volatility of the bond process $\{B_t\}_{t=0}^{\infty}$ is less than that of the government surplus process $\{\tau_t l_t - G_t\}_{t=0}^{\infty}$, provided that $\beta < 1$.*

To validate this proposition empirically, we analyze U.S. fiscal variables spanning the period 1950-2019. The federal primary surplus/deficit, compiled by the U.S. Office of Management and Budget, serves as a proxy for $\{\tau_t l_t - G_t\}$. We obtain government debt data from Feng and Santos (2023). After normalizing these data points with real GDP, we de-trend the normalized values using the H-P filter. Volatility is gauged as the ratio of standard deviation to the mean. Our findings indicate that the volatilities of the primary deficit and public debt stand at 61.5 and 19.7 respectively, corroborating Proposition 3.7.

3.2.2 General utility function

With more general utility function, the household's intra-temporal optimality condition (13) implies that λ is bounded as long as $\tau < 1$. Hence there exists an endogenous boundary for B. In the long-run, the government may choose $\mathbb{E}(q) = \beta$ to avoid permanent distortion via household's inter-temporal choice. Consequently, we have the following equation:

$$\lambda_t B_t = \sum_{i=0}^{\infty} \beta^i \lambda_{t+i} (\tau_{t+i} l_{t+i} - G_{t+i}). \tag{29}$$

The volatility of bond value process $\{\lambda_t B_t\}_{t=0}^{\infty}$ is smaller than that of the government surplus value process $\{\lambda_t(\tau_t l_t - G_t)\}_{t=0}^{\infty}$. If the utility function has a separable form, $u(c) - h(l)$, we can solve c from $\lambda = u'(c)$. If the utility function has non-separable form $U(c, 1-l)$.[4] We may still solve c from $\lambda = U_1(c, 1-l)$. From the promise-keeping condition, we have

$$\lambda = U_1(c, 1-c-g).$$

Thus,

$$d\lambda = (U_{11} - U_{12}) dc,$$

and

$$\frac{dc}{d\lambda} = \frac{1}{U_{11} - U_{12}}.$$

If $U_{12} \geq 0$, we have $\frac{dc}{d\lambda} < 0$. From the implicit function theorem, we have

$$c = c(\lambda, g).$$

To investigate the dynamics of this economy, we need some concepts.

Definition 3.8. *A Markov process $\{s_t\}_{t=0}^{\infty}$ has the Feller property if for any bounded continuous function $g(\cdot)$, $\mathbb{E}[g(s_{t+1})|s_t]$ is a bounded continuous function of s_t.*

Proposition 3.9. *Process $\{B_t\}_{t=0}^{\infty}$ is bounded.*

Proposition 3.10. *Process $\{B_t\}_{t=0}^{\infty}$ has a stationary distribution.*

The intuition behind this Proposition is that the interest rate cannot be always 0. The bounded bond process imply that the domain of the bond is endogenous. This is different from the quasi-linear case where the domain of the bond is exogenous.

Unfortunately, we could not show the ergodicity of process $\{B_t\}_{t=0}^{\infty}$. We find numerical examples which present more than one stationary distributions. Thus, the ergodicity result does not always hold for the general utility functions. Interestingly, it seems that different stationary distribution shows dramatically different properties of bond paths.

[4]Without loss of generality, we assume that the maximum labor supply is 1.

With sufficiently large initial asset position, the government may command negative interest rate to offload household's debt burden. This leads to excessive volatility in bond prices, and may amplify the business cycles.

There exists positive debt in the long run when the government has limited ability to subsidize the private sector. Bhandari et al. (2017) investigate the government's optimal policy and emphasize the mechanism through which government has a risk management role to minimize the debt risk. The interest rate or the bond price represents the risk. In the case of quasi-linear utility function, the bond price is fixed in our model.

4 Computation of Equilibria

To better comprehend the dynamics of the economy and gain further insights into equilibrium behaviors, we numerically solve the model. This necessitates the development of a computational method to characterize Ω: the set of sustainable assets/debt and promised marginal utility. This is a critical step in numerical analysis. The Ramsey problem (P-1) can be well-defined only when provided with a properly specified set Ω.

4.1 Characterization of the set Ω

To characterize the endogenous equilibrium set of feasible debt and promised shadow price of bond Ω, we draw on Abreu et al. (1990) and we apply the method developed in Feng et al. (2014) and Feng (2015). Let $\mathscr{Z}$ denote the space of aggregate shock to the economy, $\mathscr{M}$ the space of the promised shadow price of bond, and $\mathscr{B}$ the space of government budget-feasible debt B. Let $\mathscr{H}$ be the cartesian product $\mathscr{Z} \times \mathscr{B} \times \mathscr{M}$. In what follows, we define an operator $\mathbb{D}$ in $\mathscr{H}$ whose fixed point is the set of equilibrium sustainable debt and promised shadow price of bond. To facilitate the definition of $\mathbb{D}$, we begin by defining the admissibility of (B, λ) with respect to a subset $\mathbf{Y} \subset \mathscr{H}$.

Definition 4.1. *For given value function $V(z, B, \lambda)$, we say that vector $\psi = \{c, l, B_+, (\lambda_+(z_+))\}$ is admissible with respect to $\mathbf{Y}$ if ψ satisfies (P-1) for W at all (z, B, λ) and $(z_+, B_+, \lambda_+) \in \mathbf{Y}$.*

Definition 4.2. *For a given set of equilibrium sustainable debt and promised shadow price of bond $\mathbf{Y} \subset \mathscr{H}$, operator $\mathbb{D}$ is defined as*

$$\mathbb{D}^W(\mathbf{Y})(z) = \{(B, \lambda) | \exists\ \psi \text{ that is admissible w.r.t. } \mathbf{Y} \text{ at } z \text{ for } W\ \}.$$

In words, admissibility means that vector ψ guarantees existence of optimal allocations for the household and the government satisfying feasibility, while the continuation value of implied debt choice B_+, and promised shadow prices of investing in bond $(\lambda_+(z_+))$ falls within the set $\mathbf{Y}$.

Using arguments along the lines of Phelan and Stacchetti (2001), for every optimal value function $W(z, B, \lambda)$ in (P-1), we can obtain the following.

1. $\mathbb{D}(\cdot)$ is monotone and preserves compactness.

2. If $\mathbf{Y} \subseteq \mathbb{D}(\mathbf{Y})$, then $\mathbb{D}(\mathbf{Y}) \subseteq \mathbf{\Omega}$.

3. $\mathbf{\Omega}$ is compact and the largest equilibrium set of sustainable debt $\mathbf{Y}$ such that $\mathbf{Y} = \mathbb{D}(\mathbf{Y})$.

4. If we define $\mathbf{Y}_{n+1} = \mathbb{D}(\mathbf{Y}_n)$ for all $n \geq 0$, and the equilibrium set of sustainable debt $\mathbf{\Omega} \subset \mathbf{Y}_0$, then $\lim_{n\to\infty} \mathbf{Y}_n = \mathbf{\Omega}$.

4.2 Computation of equilibria via machine learning

With the above preliminaries, we start the algorithm with the following objects: (i) an initial guess for the set of sustainable debt and promised shadow price of government debt at every given z, $\Omega^{(0)}$; (ii) an initial guess for the policy $\Lambda^{(0)}(S)$, and value functions $V^{(0)}(S)$. Note that $S = (z, B, \lambda) \in \mathscr{H}$ represents the state variables.

We then solve the model by iterations. At each iteration, the algorithm takes two steps sequentially. First, for a fixed initial guess of the set of sustainable debt and promised shadow value of bond at every given z, we compute problem (P-1) through value function iteration. This step yields a new value function and policy function for the government. Next, we take as given the updated policy and value functions from the above step, and update the set of budget-feasible debt and promised shadow value of bond by implementing operator $\mathbb{D}$. These two steps must be synchronized at each iteration, and the process stops when we find convergence in all equilibrium functions and set.

The aforementioned algorithm can be computationally intense, especially regarding the numerical implementation of the operator $\mathbb{D}$. It's important to note that $\mathbb{D}$ is a set-valued operator. The methods outlined by Feng et al. (2014) and Feng (2015) involve fully discretizing the space of $\mathscr{H}$ and examining all possible elements in this space. However, the computational cost of these methods can grow exponentially with the dimension of the state space, making them less efficient. On the other hand, Judd et al. (2003) introduce a method of randomization to facilitate the convexification of the set, allowing the use of a convex hull for approximation. Despite this innovation, their method can arbitrarily enlarge the equilibrium set. Our method, in contrast, mitigates the growing computational cost issue through the utilization of random sampling. Further, by employing Deep Neural Networks (DNNs) for the approximation of the equilibrium set, we are able to improve both the efficiency and accuracy of our method, thereby overcoming the limitations found in previous approaches.

4.2.1 Deep Neural Networks

A fully connected neural network is composed of a series of fully connected layers. Every neuron in one layer is linked to all neurons in the following layer. We can understand this concept through a mathematical example as follows. Let's represent the input as $x \in \mathbb{R}^M$ and the output as $y \in \mathbb{R}^N$. To introduce a single hidden-layer fully connected network, we assume there are U_1 hidden units, denoted by $h_1^{(1)}, \ldots, h_{U_1}^{(1)}$. Then, $y_i \in \mathbb{R}$, the i-th output of y, can be computed as a function of the input x as follows:

$$h_j^{(1)} = f_j^{(1)}(x) = \sigma\left(b_j^{(1)} + \sum_{m=1}^{M} w_{j,m}^{(1)} x_m\right),\ j = 1, \ldots, U_1. \tag{30}$$

$$y_i = \sum_{j=1}^{U_1} \tilde{w}_{i,j} h_j^{(1)} + \tilde{b}_i,\ i = 1, \ldots, N. \tag{31}$$

In this equation, σ is a nonlinear activation function. Common activation functions include: $\sigma(x) = \tanh(x)$, $\sigma(x) = \max(0, x)$, and $\sigma(x) = \frac{1}{1+e^{-x}}$. The superscript 1 on $h_j, f_j, w_{j,m}$, and b_j signifies their association with a single hidden-layer function. Meanwhile, $\tilde{w}_{i,j}$ and $\tilde{b}_i$ represent the linear coefficients preceding the final output. The parameters $w_{j,m}^{(1)}$, $b_j^{(1)}$, $\tilde{w}_{i,j}$, and $\tilde{b}_j$ are all learn-able within the network.

A neuron calculates the weighted average of its inputs, and this sum is then passed through a nonlinear activation function. In the previous example, x represents the input to the artificial neurons, $f^{(1)}(x)$ represents the processing performed on the inputs, and y signifies the output through a linear combination of the neurons. It's worth noting that it's possible to stack fully connected networks directly. This can be achieved by transmitting the output of one neuron as input to the neurons of another layer, which can then perform the same computation (weighted sum of the input and transformation with the activation function). A network consisting of multiple fully connected networks is often referred to as a deep neural network (DNN). To build a two hidden-layer network, we create a set of U_1 single-layer units and treat them as input to create another set of U_2 neurons. That is, we take their linear combination and pass the result through a nonlinear activation function.[5]

$$h_j^{(2)} = f_j^{(2)}(x) = \sigma\left(b_j^{(2)} + \sum_{k=1}^{U_1} w_{j,k}^{(2)} f_k^{(1)}(x)\right),\ j = 1, \ldots, U_2. \tag{32}$$

$$y_i = \sum_{j=1}^{U_2} \tilde{w}_{i,j} h_j^{(2)} + \tilde{b}_i,\ i = 1, \ldots, N. \tag{33}$$

In (33), we retain the notation $\tilde{w}_{i,j}$ and $\tilde{b}_i$ to represent the linear coefficients associated with the final output. However, it's crucial to emphasize that these symbols signify different entities when comparing a 1 hidden-layer network to a 2 hidden-layer network. Next we can

[5]The choice of U_1 will depend on the dimension of the inputs, the number of layers chosen, and the nonlinearity of the equilibrium system.

recursively define L hidden-layer network as follows.

$$h_j^{(L)} = f_j^{(L)}(x) = \sigma\left(b_j^{(L)} + \sum_{k=1}^{U_{L-1}} w_{j,k}^{(L)} f_k^{(L-1)}(x)\right),\ j = 1, \ldots, U_L. \tag{34}$$

$$y_i = \sum_{j=1}^{U_L} \tilde{w}_{i,j} h_j^{(L)} + \tilde{b}_i,\ i = 1, \ldots, N. \tag{35}$$

It's important to note that deep neural networks are grid-free and aren't subjected to the curse of dimensionality, unlike traditional numerical techniques such as Chebyshev polynomials. Another advantage of DNNs is their "structure agnostic" nature, meaning that no special assumptions need to be made about the input. Two well-known theorems succinctly summarize the advantageous properties of DNN-based numerical approximations.

Cybenko (1989) demonstrates that neural networks, given their feed-forward architecture on the space of continuous functions, can represent a broad array of functions arbitrarily well when assigned appropriate weights. Barron (1993) establishes bounds on the error of the neural network approximation to a function, in terms of the number of hidden nodes and the smoothness of the function, independently of the dimension. This suggests that one hidden-layer neural networks can evade the curse of dimensionality.[6]

4.2.2 Outline of our numerical algorithm

One of the key steps in our algorithm is to approximate the endogenous set Ω via deep neural network.

Ω-representation: During the iteration n, we represent $\Omega^{(n)}$ with a indicator function $\chi_\Omega^{(n)} : S \to \{0, 1\}$ defined as below

$$\chi_\Omega^{(n)}(S) = \begin{cases} 1, & \text{if } S \in \Omega^{(n)} \\ 0, & \text{otherwise.} \end{cases} \tag{36}$$

We use a L_Ω-layer DNN $f_\Omega^{(L_\Omega, n)}$ to approximate $\chi_\Omega^{(n)}$, which ultimately characterize the set $\Omega^{(n)}$. Note that $f_\Omega^{(L_\Omega,n)}$ is characterized by a set of parameter values

$$\Theta_\Omega = \left\{ \left[b_{\Omega,i}^{(l)}, \left(w_{\Omega,i,j}^{(l)} \right)_{j=1}^{U_{l-1}} \right]_{i=1}^{U_l} \right\}_{l=1}^{L_\Omega} \bigcup \left\{ \left[\tilde{b}_{\Omega,i}, (\tilde{w}_{\Omega,i,j})_{j=1}^{U_{L_\Omega}} \right]_{i=1}^{N_\Omega} \right\},$$

with U_l denoting the number of nodes in layer l and $U_0 = M$. The notation $N_\Omega = 1$ indicates the dimension of the output for the function $\chi_\Omega^{(n)}$.

[6] In the Appendix, we provide one version of these two well-known theorems.

Similarly, equilibrium functions $\{\Lambda(S), V(S)\}$ will be represented by L_x-layer DNNs $f_x^{(L_x,n)}(S)$, $x \in \{\Lambda, V\}$, which are characterized by a set of parameter values:

$$\Theta_x = \left\{ \left[b_{x,i}^{(l)}, \left(w_{x,i,j}^{(l)} \right)_{j=1}^{U_{l-1}} \right]_{i=1}^{U_l} \right\}_{l=1}^{L_x} \bigcup \left\{ \left[\tilde{b}_{x,i}, (\tilde{w}_{x,i,j})_{j=1}^{U_{L_x}} \right]_{i=1}^{N_X} \right\}.$$

Likewise, N_Λ and N_V represent the dimensions of the outputs for the policy and value functions, respectively.

Update policy and value functions: During the n-th iteration of the algorithm, the endogenous set $\Omega^{(n)}$, approximated by $f_\Omega^{(L_\Omega,n)}$, is taken as given and policy and value functions are computed through a variant of reinforcement learning, c.f. Han et al. (2022).

Initially, we update the value function while considering the policy function $f_\Lambda^{(L_\Lambda,n)}(S)$ as given. The output of the policy function includes the value of λ_+ under all possible contingencies. The algorithm samples a large number N_1 of $S_{i,0} = \{z_{i,0}, B_{i,0}, \lambda_{i,0}\}$ based on $f_\Omega^{(L_\Omega,n)}$ from the space of $\mathscr{H}$. To elaborate further, we generate a set of random numbers, denoted as $\{\varsigma_i\}$, each following a uniform distribution within the interval $[0, 1]$. Each number in this set corresponds to a selection of $S_{i,0}$. We retain $S_{i,0}$ in our sample if the condition $\varsigma_i \leq f_\Omega^{(L_\Omega,n)}(S_{i,0})$ is met. Here, i is the index for the sample and 0 denotes period 0.

Next, we compute a realized utility $\hat{V}(S_{i,0})$ through a Monte Carlo simulation for T_V periods:

$$\hat{V}(S_{i,0}) = \sum_{t=0}^{T_V-1} \{\beta^t [u(c_{i,t}) - h(l_{i,t})]\} \tag{37}$$

In this equation, $c_{i,t}$, $l_{i,t}$, and all other endogenous variables can be calculated sequentially via equations (10-14), under the operation of the policy function $f_\Lambda^{(L_\Lambda,n)}$ as explained in Footnote 3.

We then update $f_V^{(L_V,n)}$ by mapping $S_{i,0}$ to $\hat{V}(S_{i,0})$ using deep learning. More specifically, we solve the following problem:

$$\min_{f_V^{(L_V,n+1)}} \mathbb{E} \left\{ \mathcal{L} \left(\hat{V}(S_{i,0}), f_V^{(L_V,n+1)}(S_{i,0}) \right) \middle| \Omega^{(n)} \right\}, \tag{38}$$

where the expectation is with respect to $\Omega^{(n)}$, and $\mathcal{L}$ denotes a loss function.[7,8] This process

[7] We establish a uniform distribution over $\Omega^{(n)}$, from which the algorithm draws N_2 random samples. The expected value, as stipulated in equation (38), is then computed as $\frac{1}{N_2}\sum_{i=1}^{N_2}\left[\hat{V}(S_{i,0}) - f_V^{(L_V,n+1)}(S_{i,0})\right]^2$.

[8] In machine learning, the loss function, denoted as $\mathcal{L}$, serves to measure the discrepancy between the predicted and actual outputs. In our context, we employ the Mean Squared Error (MSE) to determine this difference. Specifically, our loss function is expressed as:

$$\mathcal{L}\left(\hat{V}(S_{i,0}), f_V^{(L_V,n+1)}(S_{i,0})\right) = \left[\hat{V}(S_{i,0}) - f_V^{(L_V,n+1)}(S_{i,0})\right]^2.$$

results in a new neural network, $f_V^{(L_V,n+1)}$, which approximates the value function $V(S)$.

To update the policy functions, we regard the updated value function $f_V^{(L_V,n+1)}$ as a given functional form, and aim to solve the following problem:

$$\max_{f_\Lambda^{(L_\Lambda,n+1)}} \mathbb{E}\left\{ \begin{array}{c} \sum_{t=0}^{T_\Lambda-1}\left[\beta^t\left(u(c_{i,t})-h(l_{i,t})-\kappa\left(1-f_\Omega^{(L_\Omega,n)}\left(S_{i,t}\right)\right)\right)\right] \\ +\beta^{T_\Lambda} f_V^{(L_V,n+1)}\left(S_{i,T_V}\right) \end{array} \middle| \Omega^{(n)} \right\} \tag{39}$$

In this equation, $c_{i,t}$, $l_{i,t}$, and all other endogenous variables are determined as previously explained. The calculation involves an expectation with respect to $\Omega^{(n)}$, as outlined in Footnote 7. This expectation is estimated by generating N_3 samples from within the space $\Omega^{(n)}$ and then determining the average value of the enclosed term. Importantly, we introduce a term $\kappa\left(1-f_\Omega^{(L_\Omega,n)}\left(S_{i,t}\right)\right)$ into the objective function to account for the constraint (15). Here, κ is a weight parameter.

Update the endogenous set $\Omega^{(n)}$: With the updated policy and value functions $f_\Lambda^{(L_\Lambda,n+1)}\left(S;\Omega^{(n)}\right)$, $f_V^{(L_V,n+1)}\left(S;\Omega^{(n)}\right)$ in hand, we proceed to compute the endogenous set $\Omega^{(n+1)}$. The algorithm does this by sampling a large number of S_i from $\mathscr{H}$ based on $f_\Omega^{(L_\Omega,n)}$. For each S_i, we calculate the values of $\xi_i = \{c_i, l_i, \tau_i, z_{i+}, B_{i+}, (\lambda_{i+}), V(S_i)\}$ based on $f_\Lambda^{(L_\Lambda,n+1)}\left(S;\Omega^{(n)}\right)$ and $f_V^{(L_V,n+1)}\left(S;\Omega^{(n)}\right)$.

We set $\chi^{\Omega,(n+1)}(S_i)$ equal to 1 if ξ_i is admissible, otherwise $\chi^{\Omega,(n+1)}(S_i)$ equals 0. Here, ξ_i is considered admissible with respect to $\Omega^{(n)}$ if it satisfies (P-1) for $f_V^{(L_V,n+1)}\left(S;\Omega^{(n)}\right)$ at the given S_i. Numerically, the values (λ_{i+}) are calculated using the function $f_\Lambda^{(L_\Lambda,n+1)}\left(S_i;\Omega^{(n)}\right)$, while the set of endogenous variables $\{c_i, l_i, \tau_i, B_{i+}\}$ is determined through equations (10-14), as explained in Footnote 3. The admissibility requirement dictates that the resulting state $S_{i+} = (z_{i+}, B_{i+}, \lambda_{i+})$ must lie within the set $\Omega^{(n)}$. This condition is met if the value of $f_\Omega^{(L_\Omega,n)}(S_{i+})$ equals 1. Additionally, the value of $V(S_i)$ must be finite and is given by $V(S_i) = (u(c_{i,t}) - h(l_{i,t})) + \beta\mathbb{E}\left[f_V^{(L_V,n+1)}(S_{i+})\right]$.

We then approximate $\chi^{\Omega,(n+1)}$ using a DNN $f_\Omega^{(L_\Omega,n+1)}$ by solving the following problem:

$$\min_{f_\Omega^{(L_\Omega,n+1)}} \mathbb{E}\left\{\mathcal{L}\left(\chi^{\Omega,(n+1)}(S_i), f_\Omega^{(L_\Omega,n+1)}(S_i)\right)\middle|\Omega^{(n)}\right\}, \tag{40}$$

where the expectation pertains to $\Omega^{(n+1)}$ and is determined by generating N_4 samples within the domain $\Omega^{(n+1)}$, as elaborated in Footnote 7, and $\mathcal{L}$ represents the loss function. It is important to note that all these neural networks undergo training through stochastic gradient descent. The algorithm repeats the above steps until convergence is achieved in $\left\{f_\Lambda^{(L_\Lambda,n)}, f_V^{(L_V,n)}, f_\Omega^{(L_\Omega,n)}\right\}$.

This formula captures the squared difference between the estimated value $\hat{V}(S_{i,0})$ and the function $f_V^{(L_V,n+1)}(S_{i,0})$, thereby offering a clear metric of the model's performance.

5 Quantitative Analysis

In this section, we will first describe the parameterizations for the baseline economy. Following that, we will solve the model using the computational algorithm discussed in the previous section. We will then analyze the influence of various constraints on the behavior of the Ramsey government through numerical simulations.

5.1 Parameterizations

We consider two different preferences, the quasi-linear utility,

$$u(c) - h(l) = c - \gamma_l \frac{l^{1+\chi}}{1+\chi},$$

and the log utility,

$$u(c) - h(l) = \log(c) + \gamma_l \log(1-l).$$

We choose the parameters of the model as follows. The discount factor is set at $\beta = 0.9$, we let γ_l equals to 1.2 and 0.3 in quasi-linear and log preference respectively. We choose $\chi = 1.0$ so as to get a Frisch labor supply elasticity equal to 1.0. Throughout the experiments we studied in this section, the government expenditure is the only shock to the economy. It takes two values 0.155 (war) or 0.045 (peace) with the transition matrix $\begin{bmatrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{bmatrix}$. We list the details for parameter values in Table 1.

Table 1: Parameter values

Parameter	Values
β	0.9
γ_l	1.2, quasi-linear preference 0.3, log preference
χ	1.0
G	$\{0.045, 0.155\}$
P_G	$\begin{bmatrix} 0.8, & 0.2 \\ 0.2, & 0.8 \end{bmatrix}$

5.2 The case for quasi-linear preference

In Section 3.2, we discuss how a Ramsey government, operating in an incomplete market, must maintain sufficient assets to ensure solvency during severe economic shocks, referred to as G shocks. This strategy is akin to the precautionary savings behavior observed in households within income fluctuation models. In scenarios where lump-sum transfers are not accessible, the government faces challenges in accumulating an optimal asset level. This difficulty arises because the government lacks a mechanism to reinvest or manage the excess interest income generated from these assets during periods of favorable public expenditure shocks. Given the constraints of a closed economy, one alternative is for the government to lend this surplus to households. However, this approach is limited by the households' debt capacity, which depends significantly on their future net-tax earnings. These earnings are, in turn, heavily influenced by the prevailing tax policies at equilibrium.

In essence, in environments devoid of lump-sum transfers, the government's inherent asset/debt limits and its taxation capability are closely linked. The ability to tax doesn't just dictate the asset/debt limits; it also steers equilibrium dynamics. To elucidate the complex interplay between these policy tools, we consider three scenarios: *i) There's no ceiling on the government's taxation or asset/debt limitations; ii) While taxation is unrestricted, there's a defined asset limit; iii) There's a predefined tax rate limit, but no constraints on asset/debt.*

Importantly, our analysis leans more towards government asset constraints rather than debt. This focus stems from the Ramsey government's strategy to dampen and even out tax distortions by amplifying asset accumulation. Additionally, the Laffer curve's effect on tax collections inherently restricts the tax rate, thereby determining a maximum threshold for government debt accumulation.

With no lower bound on τ or B, Ponzi-scheme is optimal In the absence of any restriction on tax rate, the household's intra-temporal optimality condition (5) suggests that labor supply and earnings can increase without bound. Consequently, the natural debt limit of households, would become infinite. As discussed above, this will provide the government an instrument to absorb the surplus interest income from the accumulated assets when the economy hits by a more favorable public expenditure shock. More specifically, the government lends this extra income to the household, which results in more asset holding by the government and larger debt obligation by households. In equilibrium, the government achieves the first-best allocation by engineering a Ponzi-scheme from which the households keep borrowing from the government. This Ponzi-scheme mimics a lump-sum transfer by allowing the government to balance its budget while maintaining zero tax rate. This is sustainable as the unbounded tax policy will induce the households to labor for extended hours, hence any level of tax revenue required to close the Ponzi-scheme, when it is necessary in the future.

To illustrate the underlying economic dynamics, we model an economy where $\underline{B} = -1000$. Note that numerical simulations require finite values for meaningful computation, and we use -1000 to approximate the unbounded B. Figure 1 displays the simulated path. Evidently, the government initially builds up its asset holdings until the interest income generated is adequate to cover its expenses. The tax rate falls to a level near zero, and labor supply approximates the first-best level. Concurrently, private consumption varies inversely with public expenditure.

Figure 1: Simulation: no constraint on τ or B

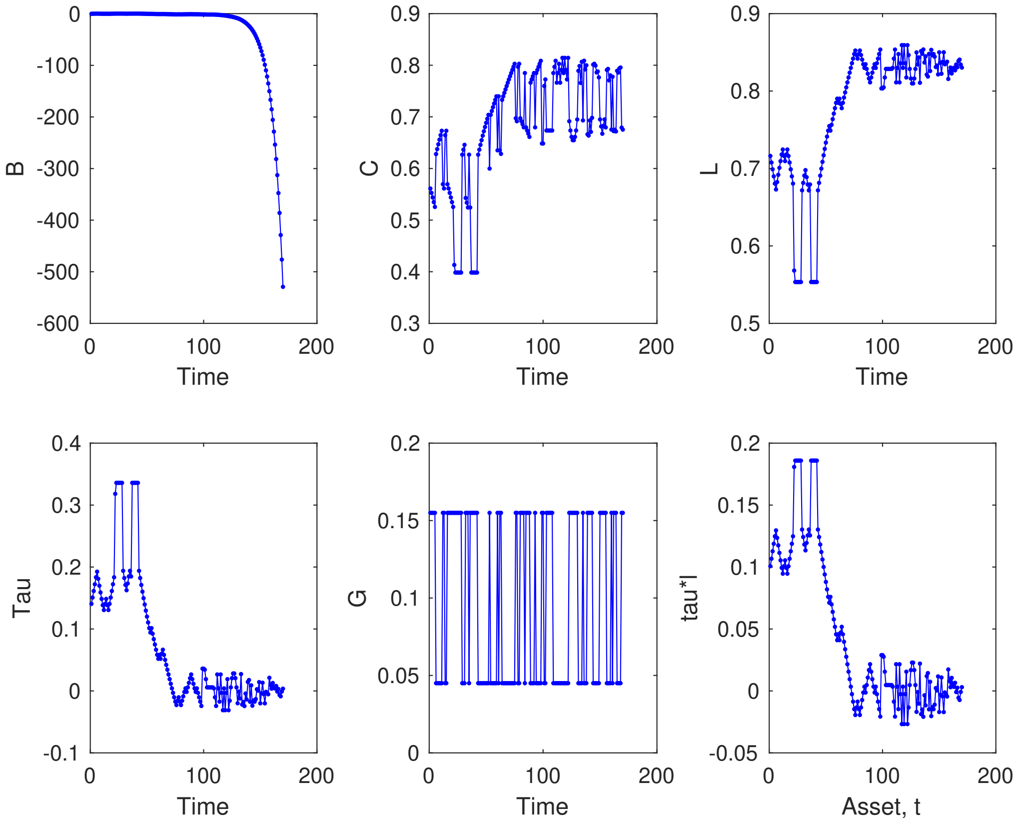

With no restriction on τ, any ad hoc asset limit is a binding constraint Subsequently, we set an arbitrary asset limit for the government. Figure 2 illustrates a scenario with an imposed ad hoc asset limit, specifically considering $\underline{B} = -1$. It's pertinent to note that these findings remain consistent across various ad hoc asset limits. Such a constraint hampers the government's ability to soak up surplus interest income by lending to households. Regardless of the asset limit level, the economy occasionally reaches this threshold, typically following a succession of favorable public expenditure shocks. This confirms the findings stated by Propositions 3.4 and 3.5.

When faced with such a shock, the government gains a primary surplus, which can be momentarily accommodated with additional asset holdings. This strategy allows the government to moderate tax rate adjustments, thereby smoothing tax distortions. However, this balancing act persists only until the asset boundary is reached. At that juncture, the government needs to recalibrate, often significantly lowering the tax rate, to ensure fiscal equilibrium.

The asset constraint also institutes a stringent minimum for tax rates. Despite the absence of explicit tax limitations, the equilibrium tax rate never dips below -7%. This phenomenon arises as the government's asset limit reflects households' ability to service their debt, contingent on their income. This income is, in turn, modulated by taxation levels.

Figure 3 displays the equilibrium value and policy functions. As discussed in Section 3.2, the value function is concave. Taxes increase with B due to the government receiving reduced interest income from bond holdings. This rise in taxes results in a decline in both labor supply and private consumption. Nonetheless, the overall tax revenue sees an uptick, aiding in balancing the budget.

Considering the policy function for bond issuance, B_+, and its close alignment with the 45° line, we focus in detail and illustrate the dynamics in Figure 4. The policy function positions itself above the 45° line during significant public spending shocks, indicating a propensity for government borrowing. In contrast, the function falls below the 45° line in response to positive public spending shocks, suggesting a government tendency to accumulate assets. Near the lower asset holding boundary, the policy function becomes flat under smaller G shocks, reflecting Proposition 3.4 that the lower bound $\underline{B}$ acts as a reflecting barrier.

Figure 2: Simulation: $\underline{B} = -1$

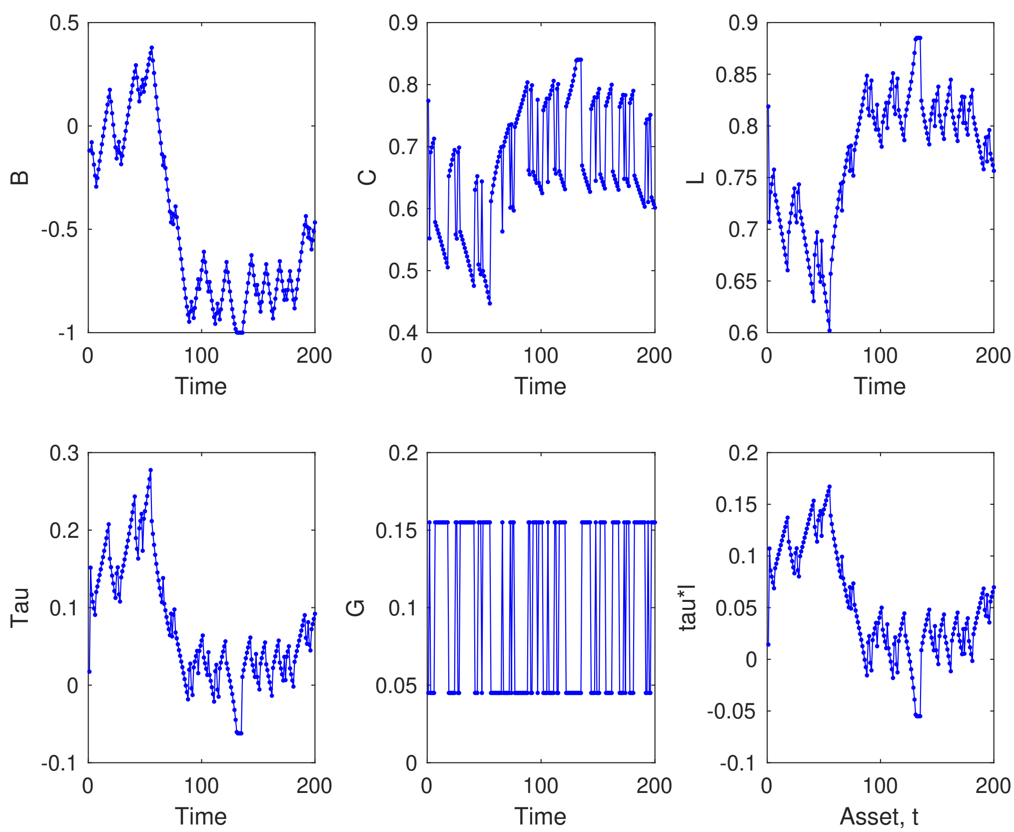

Figure 3: Equilibrium functions: $\underline{B} = -1$

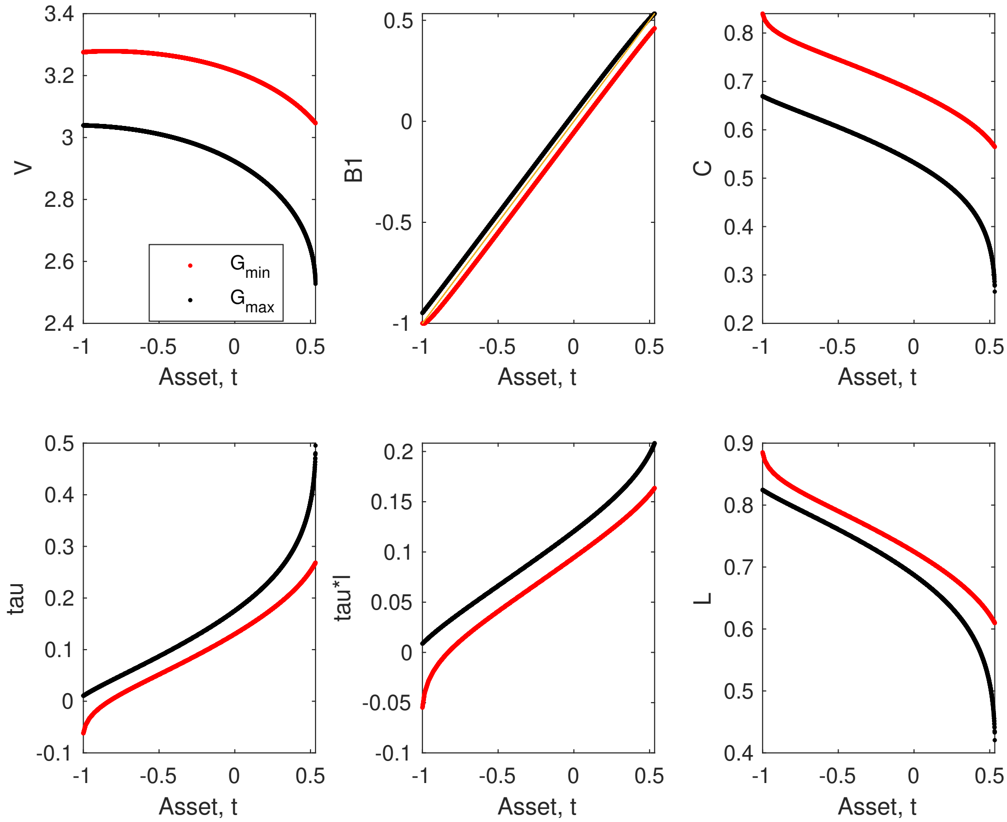

Figure 4: Policy function B_+: $\underline{B} = -1$

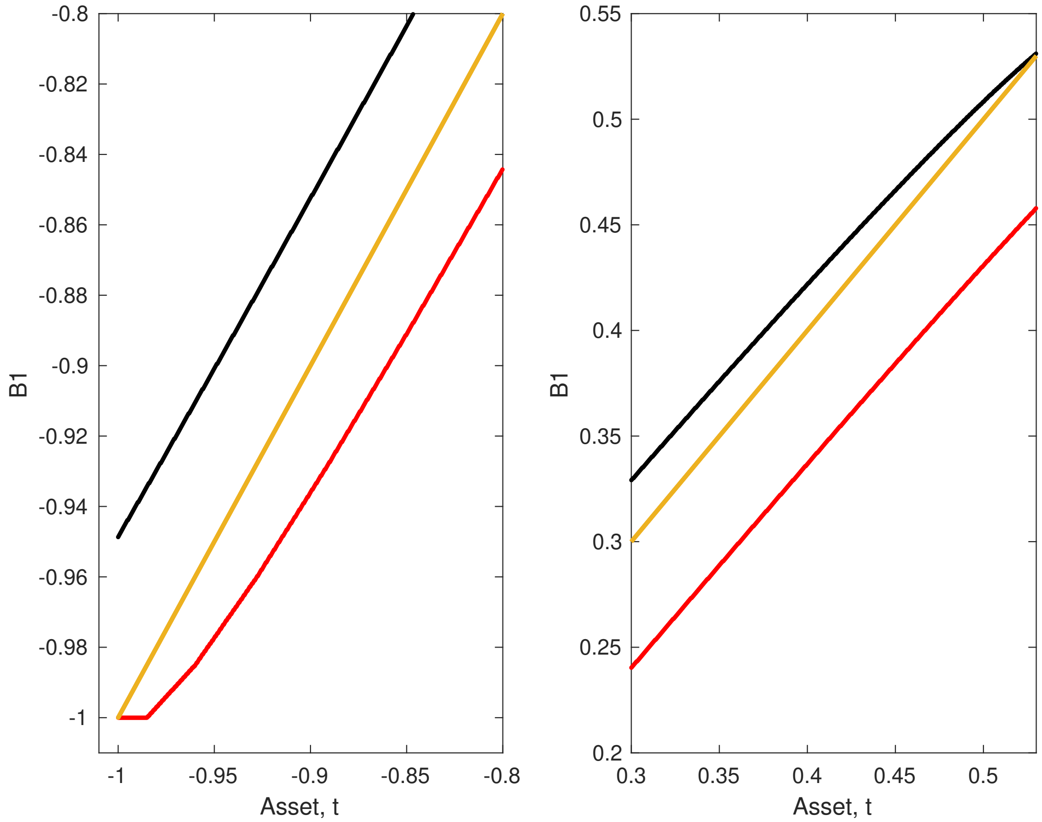

An endogenous asset/debt limit arises with a bound on τ Finally, we incorporate a tax constraint into our model. As an example, we use the condition $\tau \geq -50\%$. It's worth noting that the outcomes remain consistent regardless of the specific lower boundary set for the tax rate. This tax limit inherently shapes household income, as dictated by the household's intra-temporal optimality condition (5). Once the post-tax earnings are capped, the household's borrowing potential gets similarly restricted. In this closed economy setting, it's crucial to understand that households can only borrow from the government. This dynamic naturally sets a minimum threshold for government bond holdings, defining the so-called natural asset limit.

The related simulated data is showcased in the bottom two rows of Figure 5. Given these constraints, the government is no longer able to operate a Ponzi scheme by amassing significant assets and funding its expenditures through interest income. The natural asset and debt limits arise and are finite, with their exact values intrinsically determined by the tax rate restrictions.

Drawing a comparison to Aiyagari et al. (2002), specifically the experiment titled "incomplete markets, natural asset limit," we observe continued variations in bond issuance and tax rates even without an arbitrary asset boundary. These tests indicate that in scenarios lacking lump-sum transfers, the economy will manifest a Barro (1979) like random walk pattern in taxes and debt, albeit within certain limits.

A key point of emphasis is that with quasi-linear preference, the asset limit as discussed in Section 3.1 has an analytical representation. We have verify that the value for the lower limit on government assets generated as a fixed point of our computational algorithm matches with this implied asset limit. This congruence underscores the reliability of our numerical methodology.

Figure 5: Simulation: $\tau \geq -50\%$, natural asset limit

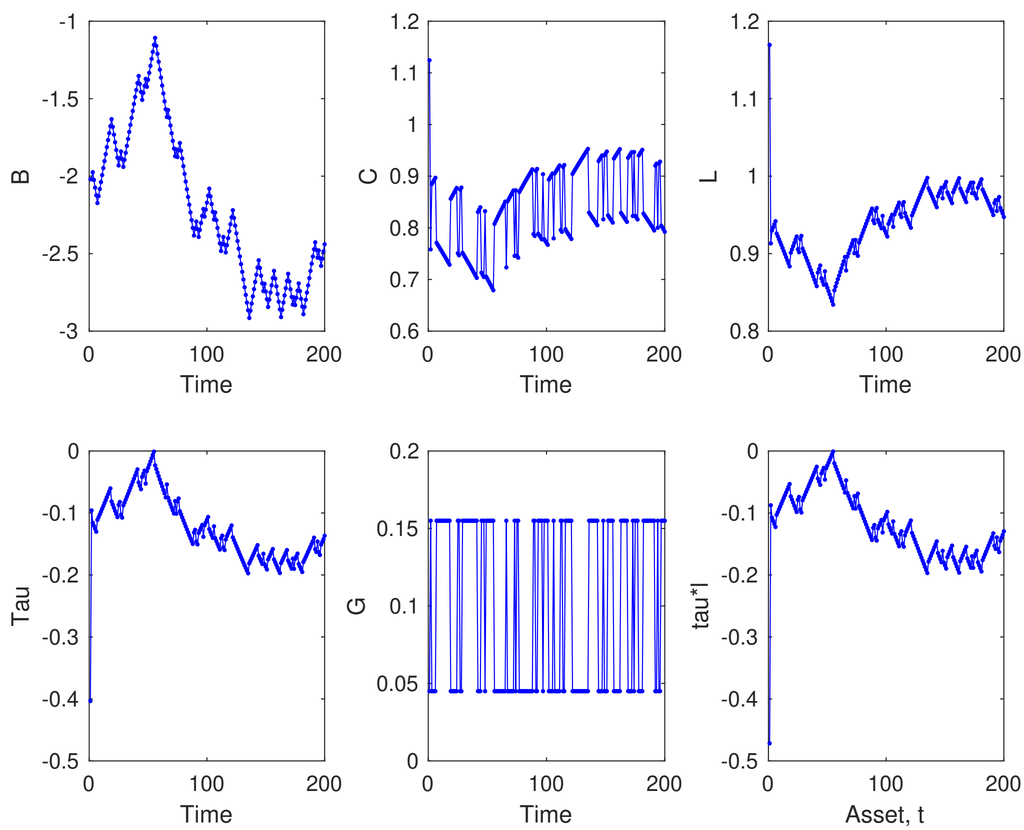

The key findings from our analysis can be summarized as follows:

- Without any restrictions on debt/assets or limits on tax rates, the economy can sustain any level of government asset. In equilibrium, the government achieves the first-best allocation by engineering a Ponzi-scheme from which the households keep borrowing from the government.

- In absence of restriction on the tax rate, any an ad hoc asset limit for the government will be a binding constraint. The government will sporadically hit the asset limit, and the economy will display Barro (1979) like random walk behavior of taxes.The asset limit will also institute a minimum for tax rates.

- When there is a lower bound limit on the tax rate, an endogenous asset and debt limit emerge. The economy exhibits a Barro (1979) like random walk behavior of taxes and debt within these boundaries.

5.3 The case for log Preference

In this section, we consider log preference. The main departure from the quasi-linear preference is that neither the marginal utility for the household nor the price for bond are constant.

Equilibrium set Figure 6 depicted below provides a numerical approximation to the set of equilibrium sustainable debt and promised shadow value of bond for an economy with logarithmic preferences. Notably, for the quasi-linear case where marginal utility from consumption is constant, this set will collapse to a straight line. As shown in Figure 6, the asset limit monotonically correlates with the promised shadow value of bond. When the Ramsey government promises a lower value of λ, it implies a higher level of consumption for the household. This can be implemented via a reduced tax rate, which encourages a higher labor supply and thereby enhances the household's capacity to accrue debt. Consequently, the sustainable asset limit monotonically increases with the value of λ.

The sustainable debt is dictated by the present value of the maximum tax revenue that the future government can amass. When the government pledges a higher shadow value of bond, it spurs a lower level of consumption and more tax revenue in the current period. However, this depresses the bond price and the present value of future tax revenue. As such, a dynamic Laffer curve emerges for debt accumulation, and the debt limit for the government exhibits a hump-shaped relationship with respect to λ.

Figure 6: Admissible set Ω

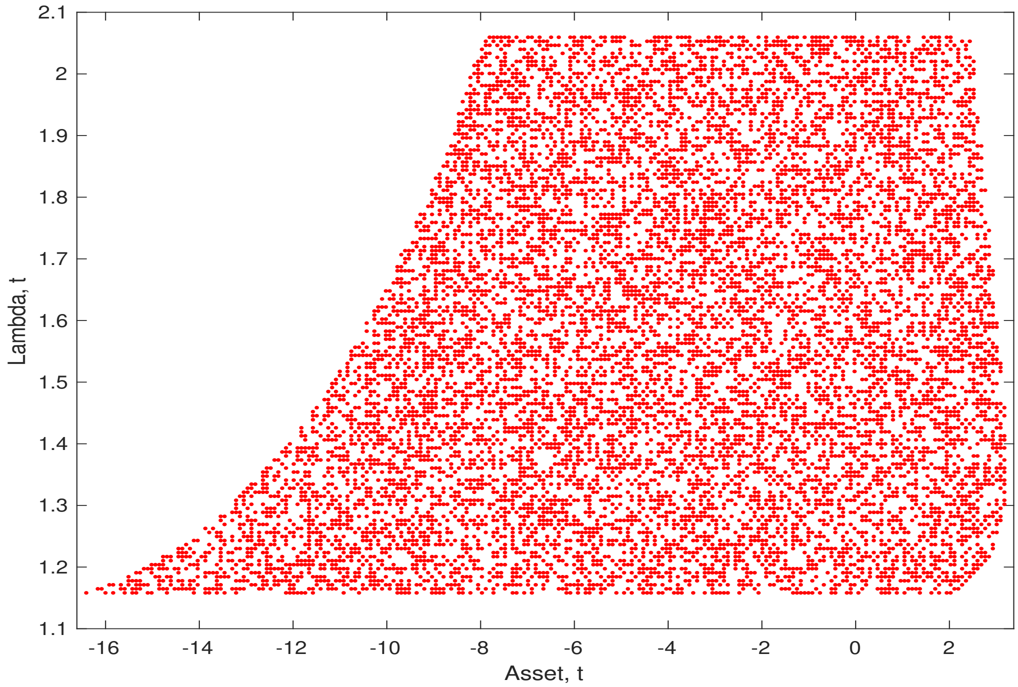

Dependence of equilibrium on initial conditions We simulate the economy considering different initial conditions, ranging from $B_0 = -29$ to $B_0 = 3$, while maintaining the same shock process for government spending across all simulations. Figure 7 and 8 display the simulated paths with $B_0 = 3$ and $B_0 = -29$, respectively. All statistical data is listed in Table 2. We observe that at least three stationary distributions exist, each indexed by the value of the government's initial asset/debt.

When the government starts with some debt or when the initial asset holding is relatively small, the economy converges to a state where both asset and tax rate experience minor fluctuations around their long-term means. The mean value of government assets as a ratio of GDP reaches 344.7%, while the average tax rate settles to -21.9% in the long run. Their respective standard deviations are 0.15 and 0.019. The bond prices converge to 0.901, almost equal to the discount factor value β.

Conversely, when the government begins with a sufficiently large asset, both bond issuance and tax rates exhibit significant volatility throughout the simulated episodes. The government amasses an asset war chest close to 900% of total GDP. The average tax rate plunges to -73.4%, triggering longer work hours and a larger total output. The average bond prices rise slightly but become more volatile, occasionally exceeding 1. In the third case, where the government begins with an asset position intermediate to the previous two scenarios, the economy behaves similarly to the case with large initial assets.

Across all simulations, both tax rate and bond prices show a negative correlation with government spending shock. As the government is initially endowed with more assets, these correlations decrease. Furthermore, the standard deviation of consumption decreases, suggesting that larger asset holdings enable the government to rely less on tax revenues to fund its spending.

Figure 7: Dependence on initial condition: $B_0 = 3$

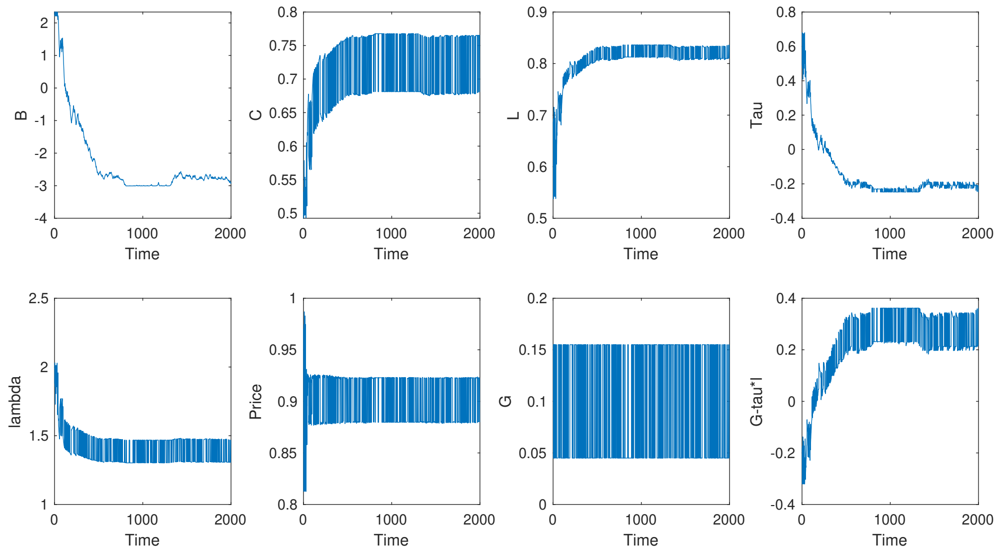

Figure 8: Dependence on initial condition: $B_0 = -29$

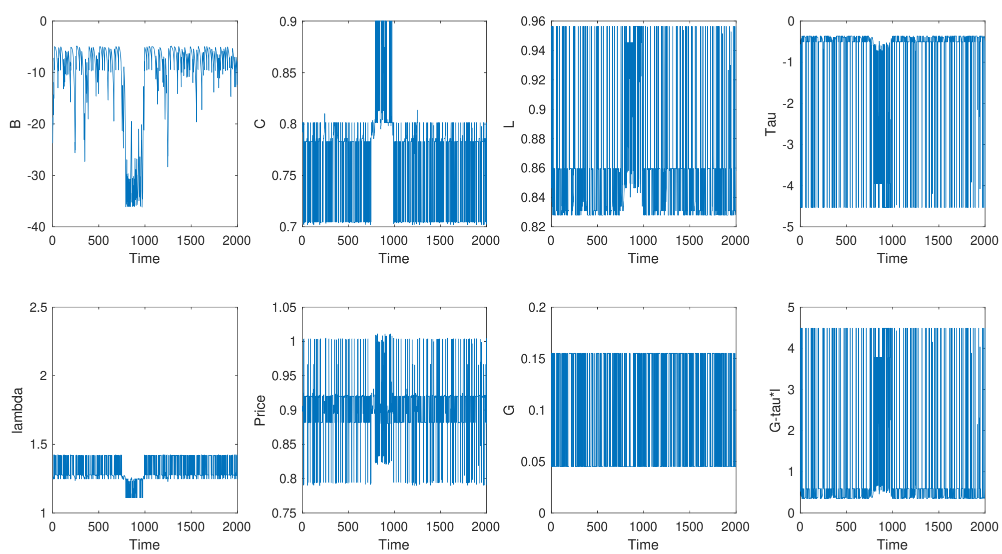

Table 2: Dependence on the initial condition

	$B_0 = 2.9$	$B_0 = 0$	$B_0 = -5.0$	$B_0 = -29$
mean($\frac{B}{Y}$)	-3.4465	-3.4465	-10.3970	-8.7979
mean(τ)	-0.2186	-0.2186	-0.8927	-0.7337
mean(Y)	0.8222	0.8222	0.8562	0.8520
mean(λ)	1.3910	1.3910	1.3280	1.3352
mean(q)	0.9010	0.9010	0.9013	0.9018
std.dev($\frac{B}{Y}$)	0.1513	0.1513	5.8281	3.1964
std.dev(τ)	0.0187	0.0187	1.2702	1.0419
std.dev(λ)	0.0827	0.0827	0.0759	0.0734
std.dev(q)	0.0214	0.0214	0.0370	0.0380
cov(τ, g)	-0.5629	-0.5629	-0.3624	-0.3347
cov(q, B)	-0.0321	-0.0321	-0.0458	-0.0860
cov(q, τ)	0.5615	0.5615	-0.3441	-0.4331
cov(q, G)	-0.9999	-0.9999	-0.4395	-0.4623
coef(τ, τ_+)	0.9992	0.9992	0.6539	0.5153

Negative interest rate In the previous subsection, we found that bond prices occasionally exceed 1 and bond issuance becomes highly volatile when the government is endowed with a sufficiently large asset position. To better understand this, we can reformulate the government's budget constraint as

$$B_t = q_t B_{t+1} + \tau_t l_t - G_t.$$

By successively substituting the government's budget constraint, we get

$$B_t = \sum_{i=0}^{\infty} \left[\left(\prod_{j=0}^{i-1} q_{t+j} \right) (\tau_{t+i} l_{t+i} - G_{t+i}) \right],$$

assuming $\sum_{j=0}^{-1} q_{t+j} = 1$.

The high volatility of $\{B_t\}_{t=0}^{\infty}$ arises from instances where q_{t+j} is occasionally greater than 1, i.e.,

$$\Pr(q_{t+j} > 1 \text{ for some } j > 0) > 0.$$

The condition $q > 1$ implies negative interest rates on government bonds. The negative interest rate triggers the explosive episode along the government bond path, which contrasts with the Martingale convergence scenario.

This argument's intuition aligns with the excess volatility of asset prices, as highlighted by Shiller and Tirole. The price of an asset is the sum of future dividends, discounted. Analogously, the government debt is equal to the sum of discounted future surpluses.

The economic implication is that large government asset limits may amplify business cycles by exacerbating the volatility of bond prices. The extensive fluctuations of government assets are a result of negative interest rates in equilibrium.

Lemma 3 of Aiyagari et al. (2002) asserts that allocations in incomplete markets cannot converge to allocations in complete markets. Our findings illustrate that incomplete markets amplify the volatility of business cycles, as complete markets cannot exhibit such large fluctuations.

Figure 9: Negative Interest Rate

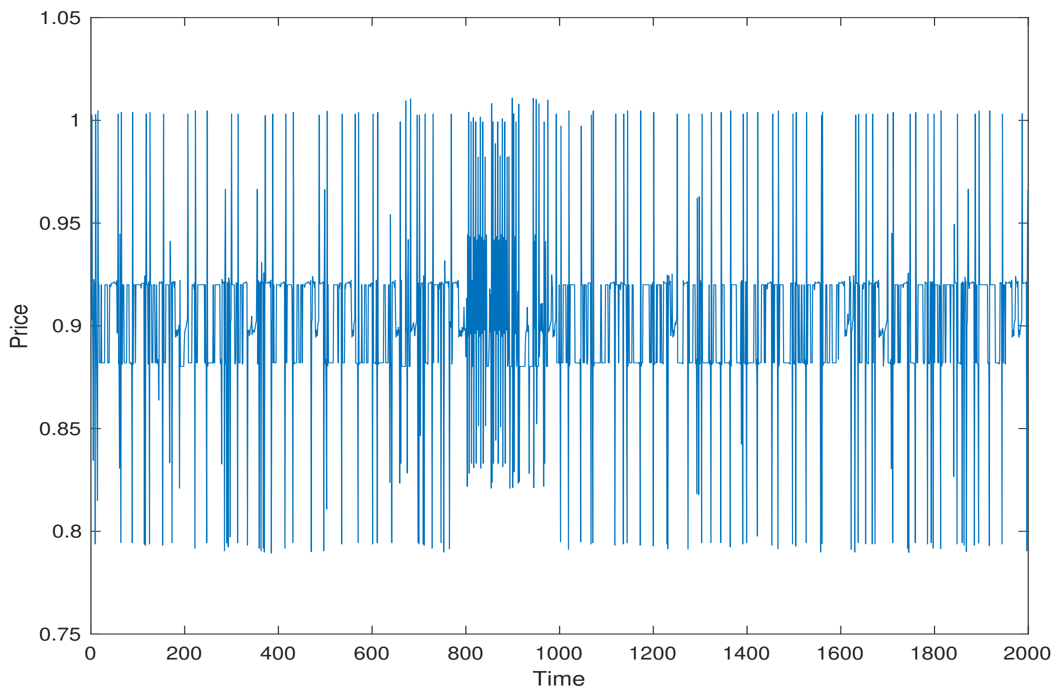

Constraints on taxation In this subsection, we explore the effects of constraints on taxation. The most salient impacts of such restrictions are observable in the asset/debt to GDP ratio and the behavior of taxation. When the government loses the capacity to impose negative taxes, the average tax rate shifts to a positive value. As a result, the government's asset position transforms into net debt. Moreover, the correlation between the tax rate and government spending shock transitions from negative to positive. Consequently, the government must rely on tax revenues to finance public expenditures.

Table 3: Impacts of constraints on τ

	$\tau_t \geq -100\%$	$\tau_t \geq 0\%$
mean($\frac{B}{Y}$)	-5.2528	0.0534
mean(τ)	-0.3983	0.1376
mean(Y)	0.8408	0.7646
mean(λ)	1.3556	1.5084
mean(q)	0.9008	0.9030
std.dev($\frac{B}{Y}$)	0.2898	0.8903
std.dev(τ)	0.0488	0.1109
std.dev(λ)	0.0761	0.1325
std.dev(q)	0.0201	0.0241
cov(τ, g)	-0.8216	0.3277
cov(q, B)	-0.1162	-0.1349
cov(q, τ)	0.8167	-0.3751
cov(q, G)	-0.9994	-0.8889
coef(τ, τ_+)	0.9960	0.9926

6 Conclusion

In this study, we explore optimal taxation, government asset management, and price volatility in an economy with incomplete markets. With a quasi-linear preference, any finite asset limit is sustainable, and in the absence of taxation constraints, the Ramsey government would resort to a Ponzi scheme. When additional limits are imposed on taxes or assets, the asset accumulation process assumes an ergodic nature even when government expenditure follows a Markov process.

Our analysis reveals that the well-documented Martingale convergence result in literature is a consequence of the government's weak precautionary savings motive. However, with a sufficiently robust precautionary savings motive from the government, the ergodic characteristic of the asset accumulation process is reinstated.

Our findings underscore the intrinsic correlation between optimal taxation and government debt/asset management, along with their considerable impact on asset pricing. This study delves into the intricate interplay between government policy, household finance, and economic stability, providing a comprehensive overview of the development and implementation of fiscal strategies within a broader economic framework.

Moreover, we introduce a machine-learning-based computational method to solve for the equilibrium set of endogenous state variables and corresponding policy and value functions. This approach is not only intuitive to implement but also adaptable to higher-dimensional problems.

References

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.

Aiyagari, S. R., A. Marcet, T. J. Sargent, and J. Seppala (2002): "Optimal Taxation without State-Contingent Debt," *Journal of Political Economy*, 110, 1220–1254.

Barro, R. J. (1979): "On the Determination of the Public Debt," *Journal of Political Economy*, 87, 940–971.

Barron, A. (1993): "Universal approximation bounds for superpositions of a sigmoidal function," *IEEE Transactions on Information Theory*, 39, 930–945.

Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent (2016): "Fiscal Policy and Debt Management with Incomplete Markets*," *The Quarterly Journal of Economics*, 132, 617–663.

——— (2017): "Fiscal Policy and Debt Management with Incomplete Markets," *The Quarterly Journal of Economics*, 132, 617–663.

Billingsley, P. (2012): *Probability and Measure*, Wiley.

Chen, X. and H. White (1998): "Nonparametric Adaptive Learning with Feedback," *Journal of Economic Theory*, 82, 190–222.

Cho, I.-K., N. Williams, and T. J. Sargent (2002): "Escaping Nash Inflation," *The Review of Economic Studies*, 69, 1–40.

Cybenko, G. (1989): "Approximation by superpositions of a sigmoidal function," *Mathematics of Control, Signals and Systems*, 2, 303–314.

Feng, Z. (2015): "Time Consistent Optimal Fiscal Policy Over the Business Cycle," *Quantitative Economics*, 6, 189–221.

Feng, Z., J. Miao, A. Peralta-Alva, and M. S. Santos (2014): "Numerical Simulation of Nonoptimal Dynamic Equilibrium Models," *International Economic Review*, 55, 83–110.

Feng, Z. and M. Santos (2023): "Cyclical Properties and Welfare of Time-Consistent Fiscal Policies," *mimeo.*

Foss, S., V. Shneer, J. P. Thomas, and T. Worrall (2018): "Stochastic stability of monotone economies in regenerative environments," *Journal of Economic Theory*, 173, 334–360.

Han, J., Y. Yang, and W. E (2022): "DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks," *mimeo.*

Jiang, W., T. J. Sargent, N. Wang, and J. Yang (2022): “A p Theory of Government Debt and Taxes,” *mimeo*.

Judd, K. L., S. Yeltekin, and J. Conklin (2003): “Computing Supergame Equilibria,” *Econometrica*, 71, 1239–1254.

Kydland, F. E. and E. C. Prescott (1980): “Dynamic Optimal Taxation, Rational Expectations and Optimal Control,” *Journal of Economic Dynamics and Control*, 2, 79–91.

Phelan, C. and E. Stacchetti (2001): “Sequential Equilibria in a Ramsey Tax Model,” *Econometrica*, 69, 1491–1518.

Renner, P. and S. Scheidegger (2020): “Machine learning for dynamic incentive problems,” *mimeo*.

Stokey, N. L., R. E. Lucas, and E. C. Prescott (1989): *Recursive Methods in Economic Dynamics*, Harvard University Press.

A Appendix

A.1 Proofs, the quasi-linear preferences

Proof of Prop. 3.6:

Proof. The process $\{B_t\}_{t=0}^{\infty}$ is generated by

$$B_{t+1} = B_+(B_t, G_t), \quad t = 0, 1, \cdots$$

We know that $B_+(B, G)$ is increasing in B.

We assume that $G_t \in \mathcal{G} \equiv \{G^1, G^2, \cdots, G^n\}$, with $0 < G^1 < G^2 < \cdots < G^n$. $\sum_{G_+} \pi(G_+|G) = 1$ for all $G \in \mathcal{G}$ and $\pi(G_+|G) > 0$ for all $(G, G_+) \in \mathcal{G} \times \mathcal{G}$.

By Corollary 1 of Foss et al. (2018) we know that the distribution of B_t converges in the uniform metric to a unique stationary distribution. □

The process $\{G_t\}_{t=0}^{\infty}$ is regenerative, since it follows a finite-state Markov chain with $\pi(G_+|G) > 0$, for all $(G, G_+) \in \mathcal{G} \times \mathcal{G}$.

We observe that debt process $\{B_t\}_{t=0}^{\infty}$ is mixing. Starting from the upper bound of the state space, the debt process has a positive probability to visiting the lower bound in finite periods. This can be achieved by keeping drawing the bad shocks of government expenditures. Starting from the lower bound of the state space, the debt process has a positive probability to surpass some positive level in finite periods. This can be achieved by keeping drawing the good shocks of government expenditures.

Proof of Prop. 3.10:

Proof. By Theorem 3.6 (Theorem of the Maximum) posited by Stokey and Lucas (1989), we know that $B_+(B, \lambda, G)$ is a continous function of B. Thus, we know that process $\{B_t\}_{t=0}^{\infty}$ has the Feller property. We also know that process $\{B_t\}_{t=0}^{\infty}$ is bounded. By Theorem 12.10 posited Stokey and Lucas (1989), we know that process $\{B_t\}_{t=0}^{\infty}$ has a stationary distribution. □

Proof of Prop. 3.1 To prove Proposition 3.1, we first show the following Lemma.

Lemma A.1. *There exists $\varepsilon > 0$ such that*

$$\Pr\left(\alpha \leq \sum_{j=0}^{\infty} \beta^j G_{t+j} \leq \alpha + \varepsilon \,\middle|\, G_t\right) \leq 1 - \varepsilon,$$

for any $\alpha \in \mathbb{R}$ and all $G_t \in \mathcal{G}$, $t \geq 0$.

Proof. Denote

$$\bar{P} = \min_{(G,G_+) \in \mathcal{G} \times \mathcal{G}} \{\pi(G_+|G)\}.$$

Choose T such that $\beta^T < \frac{1}{4}$. Let

$$\varepsilon = \min\left\{\left(\bar{P}\right)^T, \frac{\beta}{1-\beta}\frac{G^n - G^1}{4}\right\} > 0.$$

We denote

$$\bar{\alpha} = G_t + \frac{\beta}{1-\beta}\frac{G^1 + G^n}{2}.$$

Then we show this lemma in two cases.

Case (i) $\alpha > \bar{\alpha}$. Pick event $D_1 = \{G_t, G_{t+j} = G^1 \text{ for } j = 1, 2, \cdots, T\}$. On D_1 we have

$$\begin{aligned}
\sum_{j=0}^{\infty}\beta^j G_{t+j} &= G_t + \sum_{j=1}^{\infty}\beta^j G^1 + \sum_{j=T+1}^{\infty}\beta^j\left(G_{t+j-1} - G^1\right) \\
&\leq G_t + \sum_{j=2}^{\infty}\beta^j G^1 + \sum_{j=T+1}^{\infty}\beta^j\left(G^n - G^1\right) \\
&= G_t + \frac{\beta}{1-\beta}G^1 + \frac{\beta\beta^T}{1-\beta}\left(G^n - G^1\right) \\
&\leq G_t + \frac{\beta}{1-\beta}G^1 + \frac{\beta}{1-\beta}\frac{G^n - G^1}{2} \\
&= G_t + \frac{\beta}{1-\beta}\frac{G^1 + G^n}{2} \\
&= \bar{\alpha} < \alpha.
\end{aligned}$$

We know $\Pr(D_1|G_t) = \Pr\left(G_{t+j} = G^1 \text{ for } j = 1, 2, \cdots, T|G_t\right) \geq \left(\bar{P}\right)^T \geq \varepsilon$. Thus, we have $\Pr\left(\sum_{j=0}^{\infty}\beta^j G_{t+j} < \alpha|G_t\right) \geq \Pr(D_1|G_t) \geq \varepsilon$. Furthermore,

$$\begin{aligned}
& \Pr\left(\alpha \leq \sum_{j=0}^{\infty}\beta^j G_{t+j} \leq \alpha + \varepsilon \,\middle|\, G_t\right) \\
\leq\ & \Pr\left(\sum_{j=0}^{\infty}\beta^j G_{t+j} \geq \alpha \,\middle|\, G_t\right) \\
=\ & 1 - \Pr\left(\sum_{j=0}^{\infty}\beta^j G_{t+j} < \alpha \,\middle|\, G_t\right) \\
\leq\ & 1 - \varepsilon.
\end{aligned}$$

Case (ii) $\alpha \leq \bar{\alpha}$. Pick event $D_2 = \{G_t,\, G_{t+j} = G^n \text{ for } j = 1, 2, \cdots, T\}$. On D_2 we have

$$
\begin{aligned}
\sum_{j=0}^{\infty} \beta^j G_{t+j} &= G_t + \sum_{j=1}^{\infty} \beta^j G^n - \sum_{j=T+1}^{\infty} \beta^j \left(G^n - G_{t+j}\right) \\
&\geq G_t + \sum_{j=1}^{\infty} \beta^j G^n - \sum_{j=T+1}^{\infty} \beta^j \left(G^n - G^1\right) \\
&= G_t + \frac{\beta}{1-\beta} G^n - \frac{\beta^{T+1}}{1-\beta} \left(G^n - G^1\right) \\
&= G_t + \frac{\beta}{1-\beta} \frac{G^1 + G^n}{2} + \frac{\beta}{1-\beta} \frac{G^n - G^1}{2} - \frac{2\beta^{T+1}}{1-\beta} \frac{G^n - G^1}{2} \\
&= G_t + \frac{\beta}{1-\beta} \frac{G^1 + G^n}{2} + (1 - 2\beta^T) 2 \frac{\beta}{1-\beta} \frac{G^n - G^1}{4} \\
&\geq \bar{\alpha} + (1 - 2\beta^T) 2\varepsilon \\
&> \bar{\alpha} + \varepsilon \geq \alpha + \varepsilon.
\end{aligned}
$$

We know $\Pr(D_2|G_t) = \Pr\left(G_{t+j} = G^n \text{ for } j = 1, 2, \cdots, T|G_t\right) \geq \left(\bar{P}\right)^T \geq \varepsilon$. Thus, we have $\Pr\left(\sum_{j=0}^{\infty} \beta^j G_{t+j} > \alpha + \varepsilon | G_t\right) \geq \Pr(D_2|G_t) \geq \varepsilon$. Furthermore,

$$
\begin{aligned}
& \Pr\left(\alpha \leq \sum_{j=0}^{\infty} \beta^j G_{t+j} \leq \alpha + \varepsilon \,\middle|\, G_t\right) \\
\leq\ & \Pr\left(\sum_{j=0}^{\infty} \beta^j G_{t+j} \leq \alpha + \varepsilon \,\middle|\, G_t\right) \\
=\ & 1 - \Pr\left(\sum_{j=0}^{\infty} \beta^j G_{t+j} > \alpha + \varepsilon \,\middle|\, G_t\right) \\
\leq\ & 1 - \varepsilon.
\end{aligned}
$$

□

Proof of Proposition 3.1

Proof. Suppose that this is not true. We have $\Pr(|B_t| \leq M) = 1$ for some $M > 0$. Since process $\{B_t\}_{t=0}^{\infty}$ is bounded, it follows from $G_t + B_t \leq Z_t + \beta B_{t+1}$ and $Z_t \leq \bar{Z}$ that process $\{Z_t\}_{t=0}^{\infty}$ is bounded. Then there exists ψ such that we have $\Pr(Z_\infty \in [\psi, \psi + \delta]) > 0$ for any $\delta > 0$.

For any $\varepsilon > 0$, let $\eta = (1-\beta)\frac{\varepsilon}{2}$. We may choose ϕ and η, such that $\Pr(Z_\infty \in [\phi, \phi+\eta]) > 0$ and $\Pr(Z_\infty = \phi) = \Pr(Z_\infty = \phi + \eta) = 0$. Define $B = \{Z_\infty \in [\phi, \phi + \eta]\}$. Define $A_s = \{Z_s \in [\phi, \phi+\eta]\}$ and $B_s = \{Z_t \in [\phi, \phi+\eta], t \geq s\}$ for $s \geq 0$. Thus, $\lim_{s\to\infty} \Pr(A_s) = \Pr(B) > 0$ and

$\lim_{s\to\infty}\Pr(B_s) = \Pr(B) > 0$. We may choose $s < \infty$ such that $\Pr(B_s) > (1-\varepsilon)\Pr(A_s) > 0$. We have

$$\Pr\left(\sum_{j=0}^{\infty}\beta^j G_{s+j} + B_s = \sum_{j=0}^{\infty}\beta^j Z_{s+j}\,\middle|\, B_s\right) = 1.$$

Thus, we have

$$\Pr\left(\sum_{j=0}^{\infty}\beta^j G_{s+j} - \frac{1}{1-\beta}\phi + B_s = \sum_{j=0}^{\infty}\beta^j (Z_{s+j}-\phi)\,\middle|\, B_s\right) = 1,$$

and

$$\Pr\left(\left|\sum_{j=0}^{\infty}\beta^j G_{s+j} - \frac{1}{1-\beta}\phi + B_s\right| \le \frac{1}{1-\beta}\eta = \frac{\varepsilon}{2}\,\middle|\, B_s\right) = 1.$$

Let $\alpha = \frac{1}{1-\beta}\phi - B_s - \frac{\varepsilon}{2}$. We have

$$\Pr\left(\alpha \le \sum_{j=0}^{\infty}\beta^j G_{s+j} \le \alpha+\varepsilon\,\middle|\, B_s\right) = 1.$$

Since $B_s \subset A_s$ and $\Pr(B_s) > (1-\varepsilon)\Pr(A_s)$, it follows that

$$\Pr\left(\alpha \le \sum_{j=0}^{\infty}\beta^j G_{s+j} \le \alpha+\varepsilon\,\middle|\, A_s\right) > 1-\varepsilon.$$

Let $z^s = (G_0, G_1, \cdots, G_s)$. Thus, the event

$$\Pr\left(\alpha \le \sum_{j=0}^{\infty}\beta^j G_{s+j} \le \alpha+\varepsilon\,\middle|\, z^s\right) > 1-\varepsilon$$

has a positive probability since A_s is measurable with respect to z^s. Note that $\{G_t\}_{t=0}^{\infty}$ follows a Markov chain. Thus, there exists $G_s \in \mathcal{G}$ such that

$$\Pr\left(\alpha \le \sum_{j=0}^{\infty}\beta^j G_{s+j} \le \alpha+\varepsilon\,\middle|\, G_s\right) > 1-\varepsilon,$$

which contradicts Lemma 1. □

Proof of Prop. 3.2

Proof. To prove that $V(B,G)$ is differentiable at $B_0 \in (\underline{B}, \bar{B})$, note that $Z > -\infty$ implies that $G + B - \beta B_+ > -\infty$. Thus, for any B belonging to a neighborhood D of B_0, $B_+(B_0, G)$ is still feasible. Define $W(B, G)$ on D as $W(B, G) = \Gamma(G + B - \beta B_+(B_0, G)) + \beta E[V(B_+(B_0, G), G_+)|G] - G$. Thus, $W(B, G)$ is concave and differentiable in B. Since $B_+(B_0, G)$ is still feasible for $B \in D$, it follows that

$$V(B, G) = \max_{B_+ \in [\underline{B}, \bar{B}]} \{\Gamma(G + B - \beta B_+) + \beta E[V(B_+, G_+)|G]\} - G \geq W(B, G), \quad \forall B \in D,$$

with equality at B_0. Now any subgradient p of $V(B, G)$ at B_0 must satisfy

$$p(B - B_0) \geq V(B, G) - V(B_0, G) \geq W(B, G) - W(B_0, G), \quad \forall B \in D,$$

where the first inequality uses the definition of a subgradient and the second uses the fact that $V(B, G) \geq W(B, G)$ with equality at B_0. Since $W(B, G)$ is differentiable at B_0, p is unique. By Theorem 25.1 posited by Rockafellar (1970), any concave function with a unique subgradient at an interior point B_0 is differentiable at B_0. Thus, $V(B, G)$ is differentiable at B_0, Furthermore, we know that

$$V_1(B_0, G) = W_1(B_0, G) = \Gamma'(Z(B_0, G)),$$

for $B_0 \in (\underline{B}, \bar{B})$. We know that $V(B, G)$ is continuous and concave in $B \in [\underline{B}, \bar{B}]$. Thus, using Proposition 6.7.4 in Florenzano and Le Van (2001), we know that $\lim_{B \to \underline{B}} V_1(B, G) = V_1^+(\underline{B}, G)$ and $\lim_{B \to \bar{B}} V_1(B, G) = V_1^-(\bar{B}, G)$. Therefore, $V(B, G)$ is continuously differentiable in $B \in [\underline{B}, \bar{B}]$. We already know that $V_1(B, G) = \Gamma'(Z(B, G))$ for $B \in (\underline{B}, \bar{B})$. By Theorem 3.6 (Theorem of the Maximum) posited by Stokey and Lucas (1989), we know that $Z(B, G)$ is continuous in $B \in [\underline{B}, \bar{B}]$. Thus, we have $V_1(B, G) = \Gamma'(Z(B, G))$ for $B \in [\underline{B}, \bar{B}]$. □

Proof of Prop. 3.3

Proof. By Theorem 3.6 (Theorem of the Maximum) posited by Stokey and Lucas (1989), we know that $B_+(B, G)$ is a continuous function of B.

For fixed G and any $B_2 > B_1 \geq \underline{B}$, we know that either $B_+(B_1, G) = \underline{B}$ or $B_+(B_1, G) > \underline{B}$. If $B_+(B_1, G) = \underline{B}$, then $B_+(B_2, G) \geq B_+(B_1, G)$. If $B_+(B_1, G) > \underline{B}$, then we have

$$V_1(B_1, G) = E[V_1(B_+(B_1, G), G_+)|G].$$

Suppose that $B_+(B_2, G) < B_+(B_1, G)$. Then, from the Euler equqation (), we have

$$V_1(B_2, G) \geq E[V_1(B_+(B_2, G), G_+)|G] > E[V_1(B_+(B_1, G), G_+)|G] = V_1(B_1, G),$$

which contradicts the fact that $V(B, G)$ is strictly concave in B. Thus we have $B_+(B_2, G) \geq B_+(B_1, G)$. □

Proof of Prop. 3.4

Proof. Suppose that $B_+(\underline{B}, G) = \underline{B}$ for all $G \in \mathcal{G}$. Then we have

$$V_1(\underline{B}, G) \geq E[V_1(B_+(\underline{B}, G), G_+)|G] = E[V_1(\underline{B}, G_+)|G],$$

for all $G \in \mathcal{G}$. Thus we have $V_1(\underline{B}, G) = V_1(\underline{B}, G')$ for all G and G'. We have $Z(\underline{B}, G) = G + (1-\beta)\underline{B}$. This is impossible for $G \neq G'$. □

Proof of Prop. 3.5

Proof. Suppose that $B_+(B, G) > \underline{B}$ for $B > \underline{B}$ and all $G \in \mathcal{G}$. We have

$$\Gamma'(Z_t) \leq E_t \Gamma'(Z_{t+1}).$$

From $\Gamma'(Z) = \frac{\tau}{(1+\chi)\tau - \chi}$, we know that $\Gamma'(Z) \leq \frac{1}{1+\chi}$. Let $\Phi_t = \frac{1}{1+\chi} - \Gamma'(Z_t) \geq 0$. We know that

$$\Phi_t \geq E_t \Phi_{t+1}.$$

It follows from $\Phi_0 = \frac{1}{1+\chi} - \Gamma'(Z_0) = \frac{1}{1+\chi} - \frac{\tau_0}{(1+\chi)\tau_0 - \chi}$ that Φ_0 is finite. Thus, process $\{\Phi_t\}_{t=0}^{\infty}$ is a supermartingale. By the Superartingale Convergence Theorem, we know that there exists a random variable Φ_∞ with $\Phi_0 \geq E(\Phi_\infty)$ such that $\lim_{t\to\infty} \Phi_t = \Phi_\infty$ almost surely. Thus, we know that there exists a random variable Z_∞ such that

$$\lim_{t\to\infty} Z_t = Z_\infty \quad a.s.$$

However, Z_∞ could equal $-\infty$. Since process $\{B_t\}_{t=0}^{\infty}$ is bounded, it follows from $G_t + B_t \leq Z_t + \beta B_{t+1}$ and $Z_t \leq \bar{Z}$ that process $\{Z_t\}_{t=0}^{\infty}$ is bounded. Thus, there exists $G_s \in \mathcal{G}$ such that

$$\Pr\left(\alpha \leq \sum_{j=0}^{\infty} \beta^j G_{s+j} \leq \alpha + \varepsilon \middle| G_s\right) > 1 - \varepsilon,$$

which contradicts Lemma A.1. □

The process $\{(B_t, G_t)\}_{t=0}^{\infty}$ is regenerative.

Proof of Prop. 3.7

Proof. The volatility of bond process $\{B_t\}_{t=0}^{\infty}$ is smaller than that of the government surplus process $\{\tau_t l_t - G_t\}_{t=0}^{\infty}$.

$$B_t = \sum_{i=0}^{\infty} \left[\beta^i \cdot (\tau_{t+i} l_{t+i} - G_{t+i})\right].$$

Let $W(\cdot)$ be inequality measures defined over relative size that obey the Pigou-Dalton "principle of transfers." From Proposition 2 of Davies (1986), we know that $W(B_t) < W(\tau_t l_t - G_t)$. The fluctuation of government bonds is smaller than that of surplus. $\square$

A.1.1 Constraint on τ

The government's problem is

$$V(B,G) = \max_{\tau \in [\underline{\tau}, \bar{\tau}], B_+ \in [\underline{B}, \bar{B}]} \left\{ \left(\frac{1-\tau}{\gamma_l} \right)^{\frac{1}{\chi}} \frac{\tau + \chi}{1+\chi} + \beta E[V(B_+, G_+)|G] \right\} - G$$

subject to

$$G + B \leq (\gamma_l)^{-\frac{1}{\chi}} \tau (1-\tau)^{\frac{1}{\chi}} + \beta B_+.$$

We have

$$V(B,G) = \max_{Z \in [\underline{Z}, \bar{Z}], B_+ \in [\underline{B}, \bar{B}]} \{\Gamma(Z) + \beta E[V(B_+, G_+)|G]\} - G$$

subject to

$$G + B \leq Z + \beta B_+.$$

Since $\Gamma'(\underline{Z}) < \infty$, we know that the Euler equation of the government's problem is

$$V_1^+(B,G) = \max \left\{ \Gamma'(Z), E[V_1^+(B_+, G_+)|G] \right\}.$$

A.2 Theorems on numerical approximation using Neural Networks

Theorem A.2. *[Universal approximation theorem, Cybenko (1989), Theorem 5] If σ is sigmoidal, in the sense that $\lim_{z \to -\infty} \sigma(z) = 0$, $\lim_{z \to +\infty} \sigma(z) = 1$, then any function in the space of continuous functions $C([0,1]^d)$ can be approximated uniformly by two layer neural network functions.*

Theorem A.3. *[Breaking the curse of dimensionality, Barron (1993)] A one-layer neural network achieves integrated square errors of order $O(1/M)$, where M is the number of nodes. In comparison, for series approximations, the integrated square error is of order $O(1/(M^{2/N}))$ where N is the dimensions of the function to be approximated.*

A.3 General utility functions and negative interest rates

From the government's budget constraint, $B_t = q_t B_{t+1} + \tau_t l_t - G_t$, we have

$$B_t = \left(\prod_{j=0}^{n} q_{t+j} \right) B_{t+n+1} + \sum_{i=0}^{n} \left[\left(\prod_{j=0}^{i-1} q_{t+j} \right) (\tau_{t+i} l_{t+i} - G_{t+i}) \right],$$

for $n \geq 0$, with the assumption $\prod_{j=0}^{-1} q_{t+j} = 1$. Using $\lim_{n\to\infty} \left(\prod_{j=0}^{n} q_{t+j}\right) B_{t+n+1} = 0$, we have

$$B_t = \sum_{i=0}^{\infty} \left[\left(\prod_{j=0}^{i-1} q_{t+j} \right) (\tau_{t+i} l_{t+i} - G_{t+i}) \right].$$