Optimal Dynamic Fiscal Policy with Endogenous Debt Limits

Kenneth L. Judd*, Philipp Müller†, and Şevin Yeltekin‡

March 10, 2021

Abstract

Governments use debt to smooth revenues relative to spending. An important concern of economic policy is the debt capacity of an economy. We study this problem using the basic Aiyagari, Marcet, Sargent, and Seppälä (2002) incomplete market, dynamic fiscal policy model but with important changes. First, we assume that government spending is flexible. Second, we use a dynamic programming formulation for the government's dynamic problem. Third, we use global optimization methods which can handle possible nonconvexities and binding constraints. Fourth, we compute the endogenous debt limit without imposing any artificial limit. These new features lead to substantially different results compared to the previous literature. First, there is no tendency to accumulating a war chest large enough to allow taxation to disappear. Second, there are multiple ergodic sets in the long-run distribution of debt. Third, allowing for flexibility in government spending substantially increases an economy's debt capacity. Basically, if a country has a strong reputation for honoring its debt it can use both its taxing capacity and spending flexibility to accumulate substantial debt.

1 Introduction

With the recent growth in the level of government debt worldwide, the ability of a fiscal system to finance government expenditures has been a focal point in public debates. Central to these debates is the natural debt limit, i.e. the level of public debt that's sustainable in the long run, and the design of fiscal policy that is consistent with that limit. In much

*Hoover Institution, Stanford University & NBER, kennethjudd@mac.com
†Department of Business Administration, University of Zurich, philipp.mueller@business.uzh.ch
‡Simon Business School, University of Rochester, syelteki@simon.rochester.edu

of the earlier work on dynamic fiscal policy, the upper limit on debt is determined in an ad-hoc manner. Aiyagari et al. (2002)'s (AMSS) seminal paper on fiscal policy in incomplete markets revisited earlier work of Barro (1979) and Lucas and Stokey (1983) (LS) to study the implications on tax policy. Their aim was to use a general equilibrium model, as did LS, but allow only safe debt to make it comparable to the Barro model. Their results implied that taxes should roughly follow a random walk, as Barro argued. They also presented examples where the long-run tax rate is zero, and any spending is financed out of its asset income (i.e., government holds debt of the people). However, their approach had some weaknesses. First, it imposed an artificial limit on government debt and therefore did not address the question of an endogenous, natural debt limit. Second, it assumed, as much of the literature prior to it did, government spending to be exogenous. Third, it was based on only the first-order conditions of the government's dynamic optimization problem, ignoring the nonconvexities that arise naturally in optimal taxation models. Fourth, it assumed that government policy could be well approximated with a low-degree polynomial. We relax the assumptions on debt and spending, and use computational methods that aim for global optimality and allow for arbitrarily complex government decision rules. While we focus on the models examined in AMSS, we build a framework that can address a wider range of fiscal policy issues in a self-consistent manner. In particular, we derive the endogenous limits on debt and allow for endogenous government spending.

Our approach involves recasting the policy problem as an infinite horizon dynamic programming problem (see, e.g., Judd (1999)). The government's value function may not be concave and it can also very high curvature, particularly as debt approaches its endogenous limit. In dynamic taxation problems, the government's problem is a mathematical program with complementary constraints (MPCC). We explicitly use the MPCC formulation, which is essential in order to do the necessary global optimization analysis of the government's problem. Our MPCC approach uses the computational algorithms that were developed only in the past twenty years, and allows us to solve the problem reliably and accurately. The resulting dynamic programming problem is challenging in terms of both formulation and its use of computational resources. Using our combination of computational tools and more general economic assumptions, we re-address questions regarding optimal taxation and debt management in a more realistic setup. These tools allow us to determine debt limits implied by assumptions on the primitives of the economic environment and to assess how the level of debt affects both tax policy and general economic performance, and the time series properties of tax rates and debt levels.

Our results show that under the more general framework of endogenous government debt limits and spending has substantially different implications than earlier analyses have suggested. First, the behavior of optimal policy is, over long horizons (e.g., 1000 years), much more complex than simpler models imply. In particular, the long-run distribution of debt is multimodal, and the long-run level of debt is history-dependent. If initial debt is low enough and government spending is not hit with large shocks, then the government will accumulate a "war chest" which allows long-run tax rates to be zero. However, if, in the same model, initial debt is high and/or the government gets hit with a long series of bad spending shocks, then debt will rise to a high level and will not fall even if government remains low forever. In the second case, governments with large debt levels will avoid default by reducing spending and use taxes to finance a persistently high debt.

We examine the case of fixed government spending and find that the results are dramatically affected. In particular, we illustrate a case where if spending shocks are of moderate size (less than US historical experience) no positive level of debt is feasible. That is, if a government begins with positive debt then there is a sequence of spending shocks such that there is no feasible tax and borrowing policy to finance those expenditures. In such cases, exogenous spending assumptions imply that governments must have their endowed war chests in the beginning and cannot with probability one build up its war chest. These examples illustrate clearly that any analysis of fiscal policy that wants examine historical fiscal policy must consider making spending flexible.

The application of our methodology is not limited to optimal tax problems. Optimal macroeconomic policy problems, as well as social insurance design typically involve solving high-dimensional dynamic programming problems. Solving such problems is a complicated, but very important task, as the policy recommendations depend crucially on the accuracy of the numerical results. In much of the optimal macroeconomic policy and social insurance literature, accuracy of the numerical solutions is unclear. Additionally, most solution approaches ignore feasibility issues and impose ad-hoc limits on state variables such as government debt. An accurate approach to solving dynamic policy models requires the ability to handle the high-dimensional nature of the problems as well as the unknown, feasible state space. The methodology offered in this paper can be used for computing high-dimensional dynamic policy problems with unknown state spaces.

2 Related Fiscal Policy Literature

Many OLG models in the 1970's use dynamic optimization to analyze labor versus consumption versus capital taxation. Barro (1979) presented the original paper describing the time series process of optimal taxation. Under the assumption of a quadratic loss function to the government of tax rates, he showed that optimal tax policy will impose smoothly changing taxes in response to unanticipated shocks to government expenditures. Therefore, a shock to government expenditures—permanent or temporary—would be financed by a permanent increase in taxes. Lucas and Stokey (1983) examined similar issues in a general equilibrium model with complete state-contingent asset markets. In that world, tax rates depend on the movement of elasticities, and state-contingent assets absorb spending shocks.

AMSS tries to put Barro on a microeconomic foundation. They also find tax-smoothing but also present a case where the optimal policy accumulates government assets so that in the long run, all government expenditures can be financed out of a "war chest".

Our environment is similar to AMSS. We assume that the government can only issue one-period risk-free debt. The government can use a flat labor tax rate to extract resources from the representative agent. However, unlike AMSS, we do not assume that government spending is exogenous. Assuming exogenous spending substantially simplifies the analysis but is not a correct description of US history. The United States chose to enter WWI even though it faced no threat to its territory. WWII began with an attacks on the US Navy in both the Pacific and Atlantic, but the US and UK decided to go far beyond protecting their territories and their right to travel in international waters. We doubt that the exogenous spending assumption accurately describes any country's history.

We also do not put ad-hoc limits on the government debt, but endogenously determine the sustainable amount of debt. Debt is ultimately limited by the ability of the government to raise funds from taxation and keep expenditures under control. In our model, both taxation and expenditure control is used by the government to make credible its commitment to pay its debts.

A major focus in this paper is determining what debt policies are feasible. TB and LS focus on the properties of feasible policies. KP recognizes the importance of determining feasible policies. They give an excellent discussion of the problems of finding the feasible region, and we implement many of their ideas.

AMSS, TB, and LS formulate the optimal policy problem in similar ways. TB examines only the deterministic case and formulates the problem as an optimal control problem and analyzes the resulting Hamiltonian. LS takes a similar approach, formulating the solution as a time series system of Euler equations and constraints. AMSS uses the recursive contract method of MM, which is very similar to LS.

We explicitly use a utility function to model government spending choices. AMSS also has a utility function, but a strange one. Essentially, they assume that utility is zero if spending equals some target and is minus infinity at any other level of spending.

In a pointwise sense, their utility function is the limit of ours. Suppose that the target is gbar. Our utility function is -B(g-gbar)**2. At g=gbar, both of us get zero utility. For any other value of g, our utility diverges to -infinity as B goes to infinity. Therefore, our family of utility functions converges to their utility function pointwise. Therefore, we can mimic their fixed level of spending with a high value for B. Despite the pointwise convergence fact, our results show that the limiting solutions for the value function and policy functions do not converge to theirs as B increases. In the space of value and policy functions, there is a discontinuity as B approaches infinity. The AMSS specification of exogenous spending cannot be viewed as a tractable simplification of the more general case. That model is an unrepresentative special case that does not provide results that are robust to sensible alternatives. (This is a criticism that applies to many others also).

KP used a dynamic programming formulation of the optimal policy problem. This allowed them to discuss feasibility in a compact and clear manner. We follow the original approach of KP and express the dynamic problem of the government recursively with two endogenous state variables: debt and the private shadow price of wealth.

TB-LS-AMSS use only the first-order conditions. They all recognize that this assumes away nonconcavities that would give multiple solutions to first-order conditions. We do not rely on any concavity assumptions when we formulate the government's problem, and use global optimization methods that can handle the non-concave value functions and/or the nonconvex constraints sets that easily arise in optimal taxation problems.

3 The Economy

The economy is inhabited by a government and a continuum of identical households, all are assumed to be infinitely-lived. Consumers are endowed with one unit of time in each period, provide labor ℓ to produce market consumption goods c or public goods g. Time not spent in formal labor activities, $1-\ell$, can be spent at home, dedicated to leisure activities or home production.

3.1 Consumers

Utility for each consumer is a function of personal consumption, c, labor, ℓ, and government consumption, g. Furthermore, the utility of government consumption depends shock $z \in \mathbb{R}$ to its taste $\bar{g}(z)$ which follows a finite Markov process, with transition matrix $\pi(z'|z)$. We can handle the more general case, but we will follow the literature and assume utility is additively separable in consumption, government spending and labor and takes the form:

$$u(c, \ell, g, z) = uc(c) + u\ell(\ell) + ug(g, z). \tag{1}$$

We use the following parameterized specification for the utility function:

$$u(c,\ell,g,z) = \frac{(c+\underline{c})^{1-\sigma_1}}{1-\sigma_1} + \eta\frac{(1-\ell)^{1-\sigma_2}}{1-\sigma_2} - \theta(g-\bar{g}(z))^{\sigma_3}, \tag{2}$$

where θ, η, and $\underline{c}$ are positive parameters. Assuming $\underline{c} > 0$ implies that the revenue-maximizing tax rate is strictly less than unity. With this utility specification, marginal utility is finite when c=0. This allows for labor supply to be zero in response to high taxes and ensures the existence of a Laffer curve. The last term in the utility specification reflects the disutility households get from government spending that deviates from the taste shock z. The case where g is set (exogenously) exactly equal to z may seem to be the same as assuming $\theta = \infty$. However, as θ increases towards infinity, the government can still choose any former feasible policy, which—while being heavily penalized—is still feasible. Consumers evaluate the dynamic consumption, labor and government spending streams according to

$$E_0 \sum_{t=0}^{\infty} \beta^t \, u(c_t, \ell_t, g_t, z_t), \tag{3}$$

where $\beta \in (0,1)$ and E_0 denotes the mathematical expectation operator conditioned on time 0 information.

The only assets in the economy are non-contingent, real, risk-free one-period government bonds with net supply of 0. A bond at time t promises to deliver one unit of consumption at $t+1$, and has price $p_{b,t}$. At each time t, b_t is the payout of debt at beginning of current period and b_{t+1} is the number of bonds issued, which in turn mature at the beginning of time $t+1$. When $b > 0$, the government is in debt. When $b < 0$, the consumers are in debt to the government. Consumers pay a time-varying flat rate tax τ_t on their labor income. All consumers receive a transfer payment $tr_t \geq 0$ in each period. The budget constraint of a consumer at time t is given by

$$(c_t + p_{b,t} b_{t+1}) - (b_t + tr_t + (1-\tau_t)\ell_t) \leq 0, \tag{4}$$

For each consumer, the fiscal policy $\Phi \equiv \{\tau, tr, g\}$, and taste shock z are exogenous variables and the only endogenous state variable is the consumer's bond holdings. Consumers maximize their expected discounted payoff

$$\max_{\{b_{t+1}, c_t, \ell_t\}} E_0 \left\{ \sum_{t=0}^{\infty} \beta^t u(c_t, \ell_t, g_t, z_t) | \Phi_t \right\} \tag{5}$$

given fiscal policy Φ and taste shock z, subject to the sequence of intertemporal budget constraints given by Equation 4, and the non-negativity constraint on consumption, $c_t \geq 0$.

The labor supply, consumption, and debt decisions, ℓ, c and b, of the consumer depend not only upon the current state of the economy, z, and current fiscal policy Φ, but also upon future fiscal policy. Until the sequences of fiscal policies are specified, the current equilibrium decisions of the consumers cannot be determined. Suppose the optimal policy sequence chosen at time 0, $\{\Phi_t^0\}_{t=0}^{\infty}$, exists and is unique. The optimal policy will be time inconsistent in the sense that the policy $\{\Phi_t^0\}_{t=s}^{\infty}$ will not be optimal at time $s > 0$. The reason it is not optimal is because current equilibrium decisions of the consumer are functions of the current state, current policy decisions, and anticipated future policy actions. The time inconsistency severely complicates the computation of the optimal policy. Standard recursive methods are no longer applicable as shown, e.g., by Kydland and Prescott (1977). In what follows we outline a possible computational procedure following Kydland and Prescott (1980), and point out the difficulties involved.

To obtain restrictions imposed by the rational expectations equilibrium assumption, we formulate the Lagrangian for the consumer

$$L = E_0 \left\{ \sum_{t=0}^{\infty} \beta^t u(c_t, \ell_t, g_t, z_t) - \lambda_t (c_t + p_{b,t} b_{t+1} - b_t - tr_t - (1-\tau_t)\ell_t) + \mu_t c_t \mid \Phi_t \right\},$$

where λ_t is the multiplier on the t-period budget constraint and μ_t is the multiplier on t-period non-negativity constraint on consumption. The first-order conditions from the consumer problem are:

$$\begin{aligned}
\text{FOC}c: &\ -\lambda_t + uc'(c_t) + \mu_t = 0, && (6)\\
\text{FOC}\ell: &\ (1-\tau_t)\lambda_t + u\ell'(\ell_t) = 0, && (7)\\
\text{Euler}: &\ \beta \sum_{z_{t+1}} \lambda_{t+1}\pi(z_{t+1}|z_t) - p_{b,t}\lambda_t = 0, && (8)\\
\text{Budget}: &\ -b_t + c_t - tr_t - \ell_t(1-\tau_t) + p_{b,t}b_{t+1} = 0, && (9)\\
\text{KKTc}: &\ \mu_t c_t = 0, && (10)\\
&\ \mu_t \geq 0,\ c_t \geq 0. && (11)
\end{aligned}$$

Given our assumptions on the utility function, the consumer's problem is concave, so the first-order conditions are both necessary and sufficient.

The constraint given by Equation 8 has both the current shadow price λ_t and the next period shadow price λ_{t+1}. Future policies affect λ_{t+1}, which in turn affect consumers' choices in period t. When the government decides its policies, it must take into consideration the effect of its future policies on consumers' behavior in earlier periods.

3.2 Government

Each period t, the government collects labor tax revenue, pays off its old debt, b_t, issues new debt, b_{t+1}, spends g_t, and makes lump-sum transfers, tr_t. Its period t budget constraint is:

$$(\tau_t l_t + p_{b,t} b_{t+1}) - (b_t + tr_t + g_t) = 0, \tag{12}$$

We use b_t for payouts of debt at the beginning of period t. The government's plan at time t for debt payments at t+1 is denoted b_{t+1} and b^*_{t+1} represents the period demand for (and purchases of) real debt. In equilibrium, the two will be equal: $b_{t+1}= b^*_{t+1}$.

We assume a linear technology for consumption goods c and government goods g. With linear technology, the real wage w is normalized to 1. The economy-wide resource constraint at t is given by

$$(1 - \ell_t) + c_t + g_t = 1. \tag{13}$$

The timing of the moves is as follows. After the realization of the current spending shock z, the government makes its tax and spending decisions τ, tr, g, chooses market price for bonds p_b, the shadow price for the next period, and recommends allocations for the consumer c, ℓ, b^*. The consumers solve their own optimization problem, given the fiscal policy choice and will pick the same allocation as suggested by the government, if it's in their interest to do so. To ensure that consumers follow through with his plan, the government chooses fiscal policy and consumer allocations that are consistent with consumers' optimal choices of consumption and labor. Additionally, these policies must deliver the shadow price λ^* and debt b^* from the consumer's problem.

In the recursive formulation of the government's problem, we drop the t subscripts and use superscript + to denote next period's variables. The shock z, is an exogenous state variable, and the endogenous state variables for the government problem are b and λ. The dynamic programming problem the government solves is the following.

$$V(b, \lambda, z) = \max \; u(c, \ell, g, z) + \beta \mathbb{E}[V(b^+, \lambda^+, z^+)] \tag{14}$$

subject to its budget constraint

$$(\tau \ell + p_b b^*) - (b + tr + g) = 0,$$

aggregate resource constraint

$$\ell = c + g,$$

promise-keeping constraints for λ and b

$$\begin{aligned} \lambda^* &= \lambda \\ b^* &= b^+, \end{aligned}$$

and the first-order conditions from the household's problem,

$$\begin{aligned} -\lambda^* + uc'(c) + \mu &= 0 \\ (1-\tau)\lambda^* - u\ell'(1-\ell) &= 0 \\ \beta \sum_{z^+} \lambda^+ \pi(z^+|z) - p_b \lambda^* &= 0 \\ -b + c + tr - \ell(1-\tau) + p_b b^* &= 0 \\ \mu c &= 0, \end{aligned}$$

non-negativity constraints,

$$c, \ell, g, p_b, \lambda^*, \lambda^+, \mu, tr \geq 0,$$

and the government feasibility constraint

$$(b^+, \lambda^+) \in \Omega(z^+),$$

where $\Omega(z^+)$ is the set of (b, λ) for which there exists a policy sequence with an equilibrium in state z^+.

The constrained optimization problem defined by Equation 14 is not a standard dynamic programming problem. First and foremost, the set of feasible b and λ combinations, Ω, is unknown. Second, the constraint set is not convex. Third, the value function may not be concave. These issues present nontrivial computational challenges which we address next.

4 Computational Algorithm

Solving the dynamic policy problem of the government presents many computational challenges. Before approximating the government's value function, we must study the unknown domain. More precisely, we must identify the combinations of b and λ that are economically feasible. This feasibility issue was recognized by Kydland and Prescott (1980) who suggested an iterative procedure that embeds the dynamic programming problem of the government into a fixed point problem for finding the a priori unknown feasible region $\Omega(z)$, but did not provide an actual algorithm. They acknowledged that

> "This formulation leads to unusual constraints, however, and the problem of actually computing an optimal policy would appear quite formidable even for relatively simple parametric structures."

Determining the feasible domain is not the only computational issue. The set $\Omega(z)$ may not be rectangular or even convex, which further complicates the approximation of the value function because the value function is likely to badly behave along the "feasible" boundary. In particular, the marginal value of debt will be close to the marginal social cost of revenue, which will diverge as the revenue approaches the maximum possible. The approximation procedure for the value function must be flexible enough to handle such a possibility.

Due to the computational difficulties described above, we use *discrete state dynamic programming* to approximate the value function and policy function. It may appear natural to uniformly discretize the promised marginal utility of consumption λ and government debt b for peace and war. For utility functions satisfying the Inada conditions, however, λ tends towards infinity for $c \to 0$. Instead, we discretize the consumption c uniformly from c_{min} to 1. The grid is then non-uniformly spaced in λ with $\lambda \in [uc'(1+\underline{c}), uc'(c_{min})]$. The b domain depends on the utility function. We choose $c_{min} = 0.01$.

The log-log utility function specification allows for a closed-form solution of the optimization problem for any policy $(b^+, \{\lambda^+\}_{z+})$. Note that the $|(b^+, \{\lambda^+\}_{z+})| = 1 + |z^+|$. We exemplary derive the closed-form solution for $\underline{c} = 0.1$ and flexible g: First we focus on the non-negativity constraint on tr and assume that $tr > 0$. Solving for the unconstrained problem, τ equals

$$\tau_{tr>0} = \frac{-100 + 410\lambda - 100\bar{g}(z)\lambda - 219\lambda^2 \pm 10\sqrt{100 - 20\lambda + 200\bar{g}(z)\lambda + 3\lambda^2 - 20z\lambda^2 + 100\bar{g}(z)^2\lambda^2}}{200\lambda - 219\lambda^2 + 200z\lambda^2}. \tag{15}$$

We can discard the "−" case as it yields imaginary numbers for the utility. Next we check our assumption $tr > 0$ with

$$tr = \frac{100 - 55\lambda - 50b\lambda + 50\lambda\tau + 48b^+\mathbb{E}[\lambda^+]}{50\lambda}. \tag{16}$$

If $tr \geq 0$ is indeed non-binding, we can proceed, otherwise, set $tr = 0$. In this case, we calculate

$$\tau_{tr=0} = \frac{-100 + 410\lambda - 100\bar{g}(z)\lambda - 219\lambda^2 \pm 10\sqrt{100 - 20\lambda + 200\bar{g}(z)\lambda + 3\lambda^2 - 20\bar{g}(z)\lambda^2 + 100\bar{g}(z)^2\lambda^2}}{200\lambda - 219\lambda^2 + 200\bar{g}(z)\lambda^2}. \tag{17}$$

Using the correct τ from above, the remaining choice variables follow as

$$c = -\frac{1}{10} + \frac{1}{\lambda} \tag{18}$$

$$\ell = 1 - \frac{50 - 5\lambda - 50b\lambda + 48b^+\mathbb{E}[\lambda^+]}{100 - 5\lambda - 50b\lambda + 48b^+\mathbb{E}[\lambda^+]} \tag{19}$$

$$p = \beta\frac{\mathbb{E}[\lambda^+]}{\lambda} \tag{20}$$

$$g = \frac{1000 - 50(13 + 10b)\lambda + 55(1 + 10b)\lambda^2 + 48b^+(10 - 11\lambda)\mathbb{E}[\lambda^+]}{10\lambda(5(-20 + \lambda + 10b\lambda) - 48b^+\mathbb{E}[\lambda^+])}. \tag{21}$$

We check the inequality constraints $\tau \leq 1$ and $g, p_b, \ell \geq 0$. If any constraint is violated, the policy $(b^+, \{\lambda^+\}_{z+})$ is infeasible.

The algorithm starts with identifying the numerically feasible region $\tilde{\Omega}(z) \subseteq \Omega(z)$ by iteratively improving the k-th guess approximation of the feasible region $\tilde{\Omega}^k(z)$ until we found its fixed point, i.e., $\tilde{\Omega}^{k+1}(z) = \tilde{\Omega}^k(z)$.[1] To initialize, we set the policies to "stay where you are" and evaluate the policies. This yields an initial distribution of feasible/infeasible states $\tilde{\Omega}^k(z)$ with $k = 0$. We find $\tilde{\Omega}^{k+1}(z)$ by evaluating for each infeasible point all possibly feasible policies in the search space $(b^+, \{\lambda^+\}_{z+}) \in \tilde{\Omega}^k(z)$. While the initial search space comprises all feasible points in the state space $\tilde{\Omega}^k(z)$, it can be reduced to the "newly" feasible points after the first iteration. Note that this procedure is possible without loss of accuracy. This iterates until we have found the fixed point.

If the set of feasible points converged, the algorithm continues alternating policy improvement and iteration steps. The policy improvement searches for each point the optimal policy globally in $\tilde{\Omega}(z)$. The policy iteration iterates until the L^1-error is less than 10^{-13}. Both steps alternate until there is no policy improvement.

[1] Note that the numerically feasible region $\tilde{\Omega}(z)$ is a subset of the actual feasible region $\Omega(z)$ due to the discretization of the state space.

Case #	σ_1	σ_2	σ_3	$\underline{c}$	exogenous g	$\bar{g}$	b bounds	c bounds
1	1.0	1.0	2	0	No	{0.09, 0.27}	[-10, 12]	[0.01, 1]
2	1.0	1.0	-	0	Yes	{0.09, 0.27}	[-10, 12]	[0.01, 1]
3	1.0	1.0	2	0.1	No	{0.09, 0.27}	[-10, 12]	[0, 1]
4	1.0	1.0	-	0.1	Yes	{0.09, 0.27}	[-10, 12]	[0, 1]

Table 1: Experiment setup

5 Results

In this section, we present the computational results from a variety of parameterizations of our model. The utility function for the consumers, as mentioned before, takes the form

$$u(c, \ell, g, z) = \frac{(c + \underline{c})^{1-\sigma_1}}{1 - \sigma_1} + \eta \frac{(1 - \ell)^{1-\sigma_2}}{1 - \sigma_2} - \theta(g - z)^{\sigma_3}. \tag{22}$$

For all of the cases reported, we set $\eta = 1$ and $\theta = 100$. The discount factor β is fixed at 0.96. The transition matrix equals in all cases

$$\Pi = \begin{bmatrix} 0.9787 & 0.0213 \\ 0.3333 & 0.6667 \end{bmatrix}, \tag{23}$$

roughly following Buera and Nicolini (2004) in their calibration of Π to the US experiences in the last century: roughly a war twice a century with an average duration of 3 years. All experiments were run on an NVIDIA Tesla V100 GPU.

Table 1 reports the specifics of the four cases we compute. The bounds on debt b bounds and consumption c are the state spaces that include both the economically feasible and infeasible points. For each case, we display the endogenously determined, economically feasible/infeasible regions inside these sets. For expositional reasons, all plots are plotted in the (λ, b) domain with $\lambda = uc(c)$; the log-log cases plotted with $\log_1 0(\lambda)$. In cases 1 and 3, government spending g is a choice variable. In cases 2 and 4, government spending is constrained to equal the exogenous shocks and these cases provide a direct comparison to the results in AMSS.

5.1 Feasible Regions

Figure 1 displays the feasible/infeasible regions for all cases in Table 1. In cases 1 and 3—endogenous government spending and $\bar{c} \in \{0, 0.1\}$—the feasible region is non-convex; as λ increases, higher debt levels become infeasible. In other words, as debt levels increase, high shadow values of debt cannot be supported in equilibrium. In cases 2 and 4—exogenous government spending and $\bar{c} \in \{0, 0.1\}$—the feasible region is convex and restricted to the second quadrant; the government cannot borrow from the households. This is sensitive to the exogenous government spending specification as shown in Appendix A.1.

5.2 Simulations

The central question we want to answer is how debt evolves over time, as the economy receives low and high spending shocks. Is there a tendency to always lower debt and accumulate a war chest? If a war shock hits the economy today is there always an accumulation of debt? We now display the answers for Case 3.

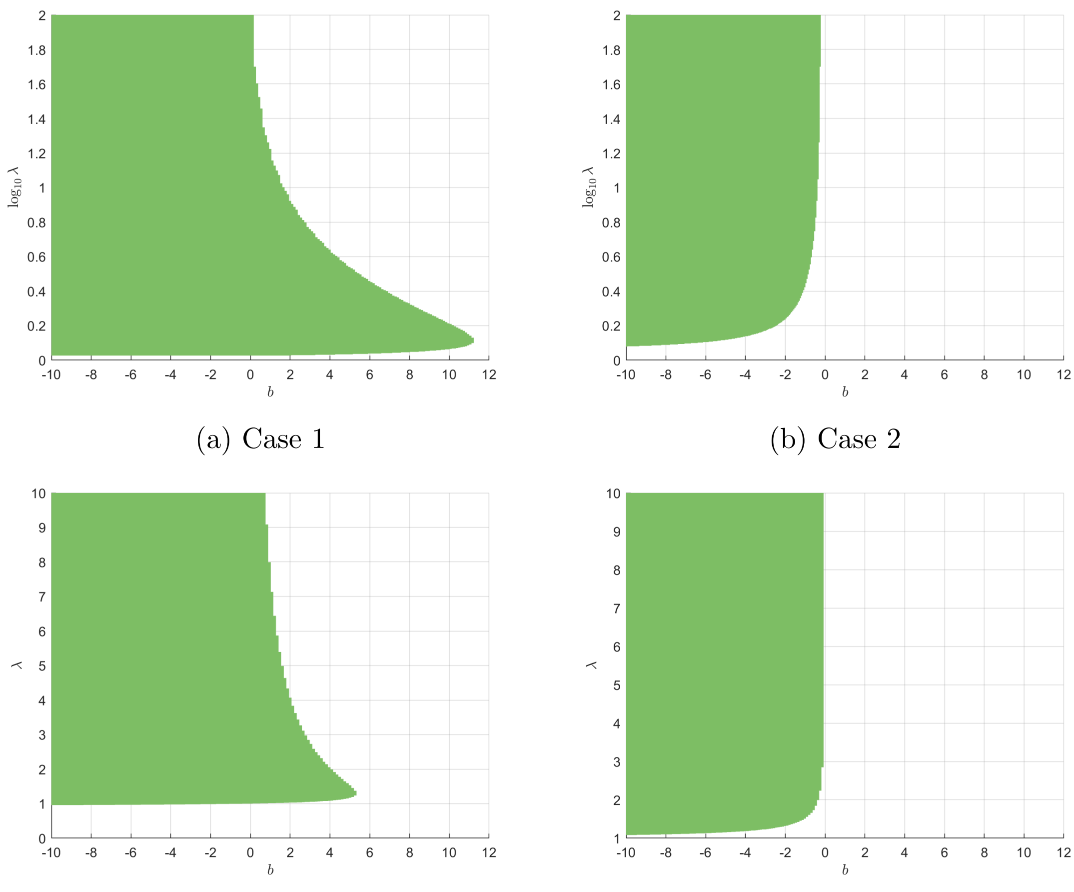

(a) Case 1 (b) Case 2

(c) Case 3 (d) Case 4

Figure 1: Numerically feasible regions $\tilde{\Omega}(z)$ for cases in Table 1.

5.2.1 Perpetual Simulations

We display the paths followed if the economy begins with $(b, \lambda) = (\text{binit}, \text{linit})$ and proceeds with no change in the expenditure state for 50 periods.

In perpetual peace in Figure 2, we see that both the b and λ paths move quickly to their long-run values. Notice that for any fixed initial λ, the long-run value of (b, λ) depends on the initial value of b. There is no drift to a war chest nor to high debt. It appears that there are many (b, λ) points that do not change as long as the state is peace.

The perpetual war case in Figure 3 is different. We find that it takes many periods for some long-run state is achieved. For most initial levels of debt, there is a gradual climb in debt. Future graphs will also show that there are declines in consumption and government expenditure. The intuition is clear. When in war, there is a 1/3 chance of peace the next period. Therefore, it is natural to use debt to finance a level of expenditure that is thought to be high relative to future expenditure. However, if the state remains war, the rise in debt has to hit a ceiling.

If b starts sufficiently low, then b is a war chest, no taxes are necessary and b will remain at a war chest level. We see that for low initial values for b. However, if b starts above the war chest level, then b will climb slowly to a limit, which appears to be roughly the same for all initial positive levels of initial b.

Not only do these results overturn the AMSS "always build a war chest" result, they may seem counter intuitive at first. Perpetual wars do not seem to wipe out war chests and perpetual peace does not seem to wipe out positive debt. The answer lies in the ability of the government to choose the level of spending, instead of being exogenously constrained to always spend the amount corresponding to the spending shock.

5.2.2 Long-run States

We next display the movement of debt and λ for a stochastic path of peace and war, generated according to the Markov transition matrix we assumed. We simulated the model for a large collection of initial (b, λ) values. We found four kinds of long-run states.

If the initial b is small, then we converge to two points in (b, λ) space, one for peace and one for war, where the level of b is the same for each. This is the war chest situation where b is constant but λ changes because consumption much respond to the change in taste for government expenditure. We do not display that simple case.

In each of the next figures, we show six point plots. The first shows the (b, λ) states visited for the first one hundred periods, as indicated by the label $1 \leq t \leq 100$. Each of the five other plots displays the points visited over some later period of time.

Perpetual peace case

8]//TableForm=

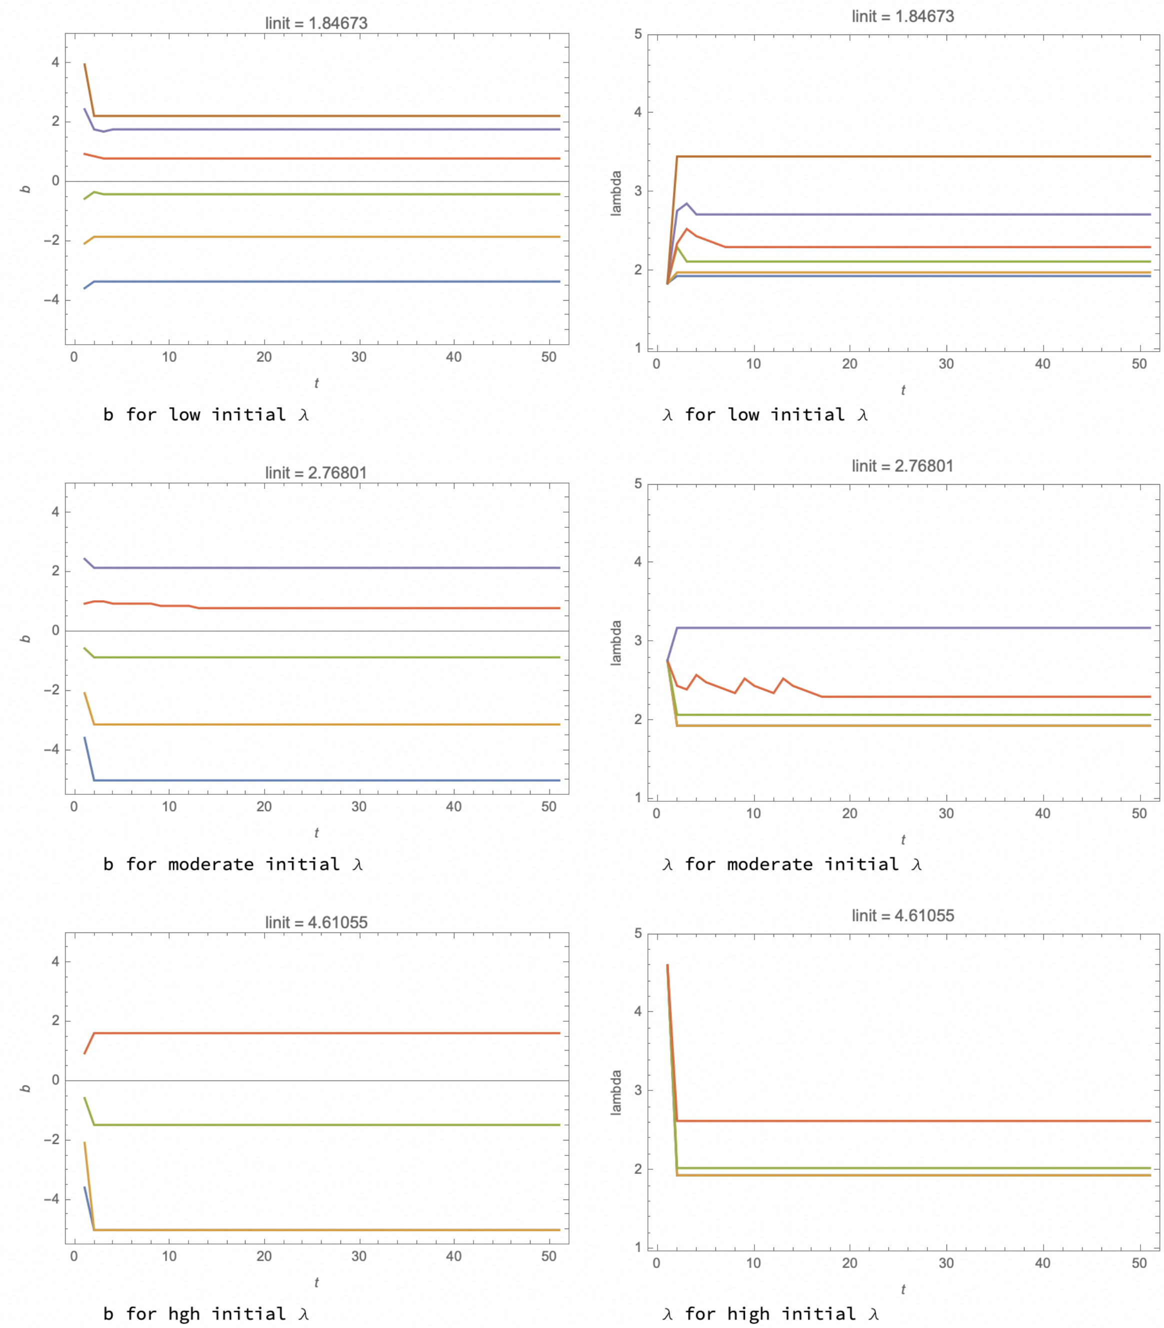

Figure 2: Perpetual Peace

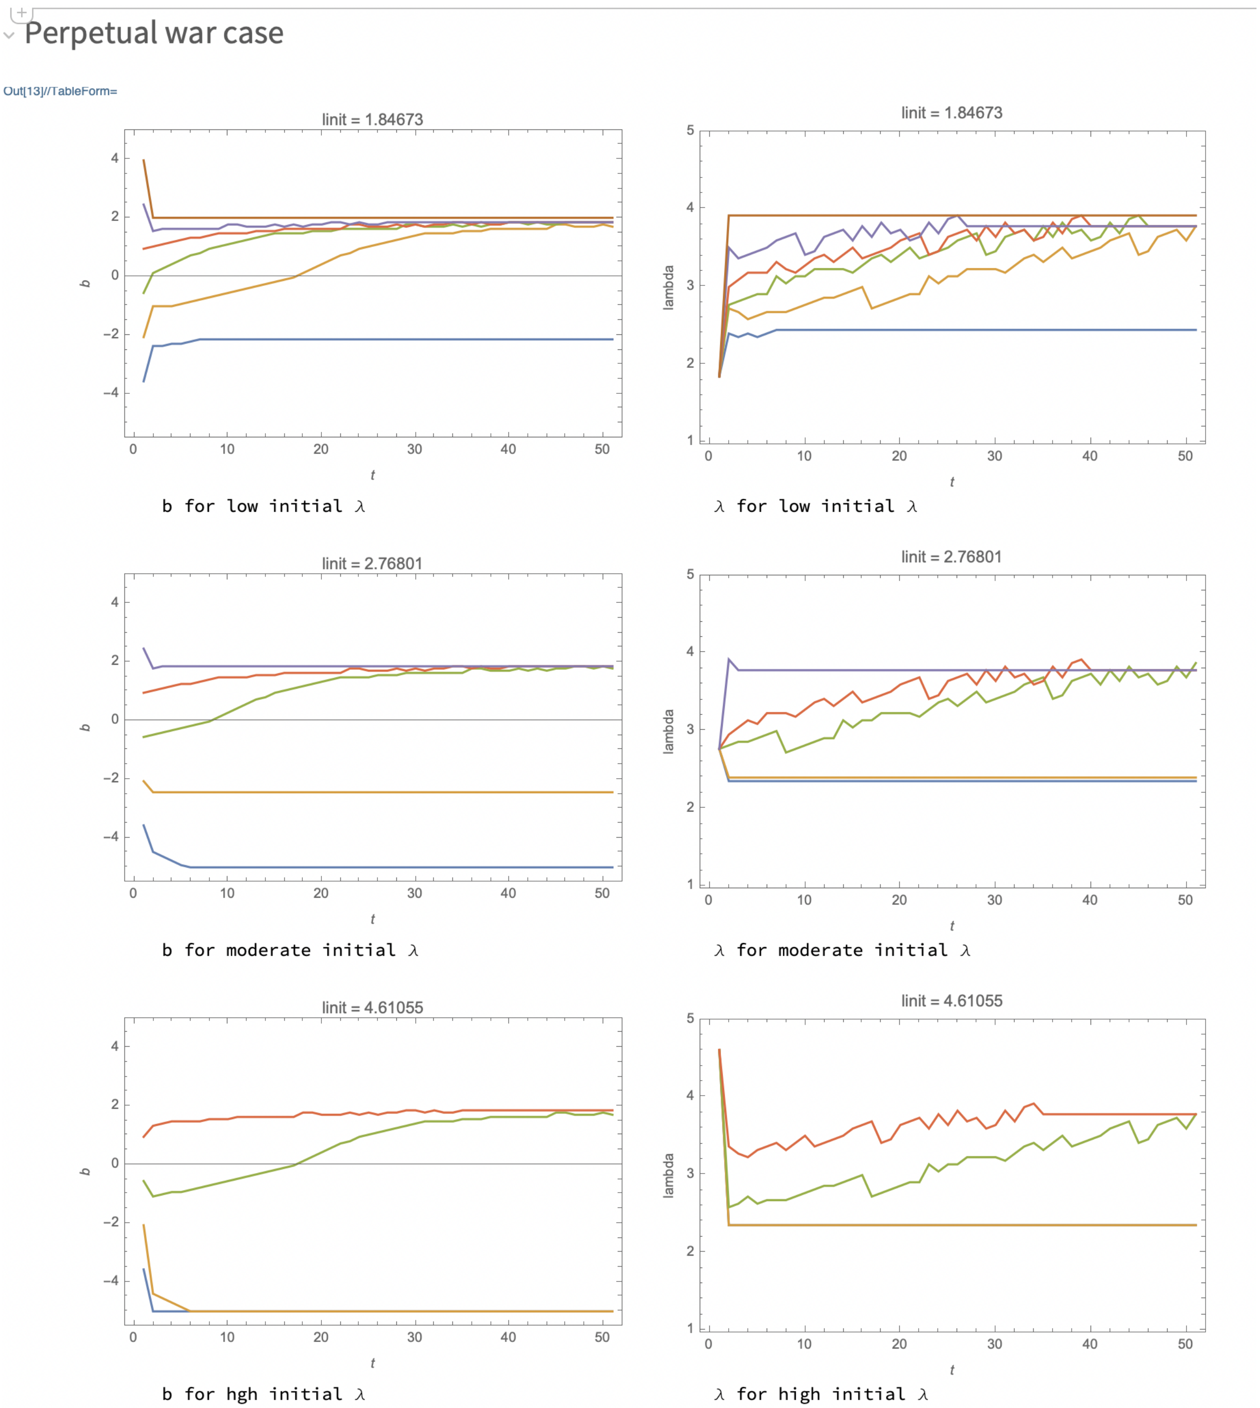

Figure 3: Perpetual Peace

Figure 4 displays a rather strange evolution. There is not much movement but the last 4250 periods are spent at just four points. This is a rather fragile result that happens for only a few initial states.

Figure 5 is typical if initial debt is positive or slightly negative. There is a slow drift towards high debt. In fact, the last 4250 periods are spent close to the boundary of the feasible state. Note that this drift is very slow. The long-run ergodic states are so far in the future that they are irrelevant for any discussion of what happens when debt is small.

Figure 6 is also a knife edge case arising for a few initial conditions. The initial debt level is too high to be a war chest, but only barely so. In this case we find that there still is a drift up in debt but the debt level increases only by 0.5 over the course of 5000 periods.

6 Conclusion

We can solve the AMSS model with high precision if we use the dynamic programming approach advocated by Kydland and Prescott (1980). When we assume exogenous spending as done in AMSS, we find that the feasible region for debt is small for historically reasonable assumptions for government spending. In fact, it is easy to find cases where no positive initial level of debt is feasible.

We, however, assume that government expenditure is flexible and chosen according to a utility function. This is a far more reasonable assumption about government spending. We find that the set of feasible debt levels is substantially larger. The reason is clear: flexibility in spending allows the government to credibly borrow funds because its creditors know that spending will be cut if necessary to honor the debt.

Our base example also displays behavior inconsistent with simple autoregressive models. While the solution is stationary because it is a Markov chain, the convergence to the stationary state is so slow that it takes longer than 5000 periods for it to be apparent.

The final point we make is that it is feasible to solve such models with high accuracy with modern algorithms, software, and hardware. As this project proceeds, we will build a large library of plots displaying results for a large number of cases. This will allow us to determine how the key parameters affect the results. We will not rely on a couple of cases to make assertions about why the results are what they are. With this library (which will be posted in the web), we will be able to explore alternative hypotheses based on the data provided by the library of simulations.

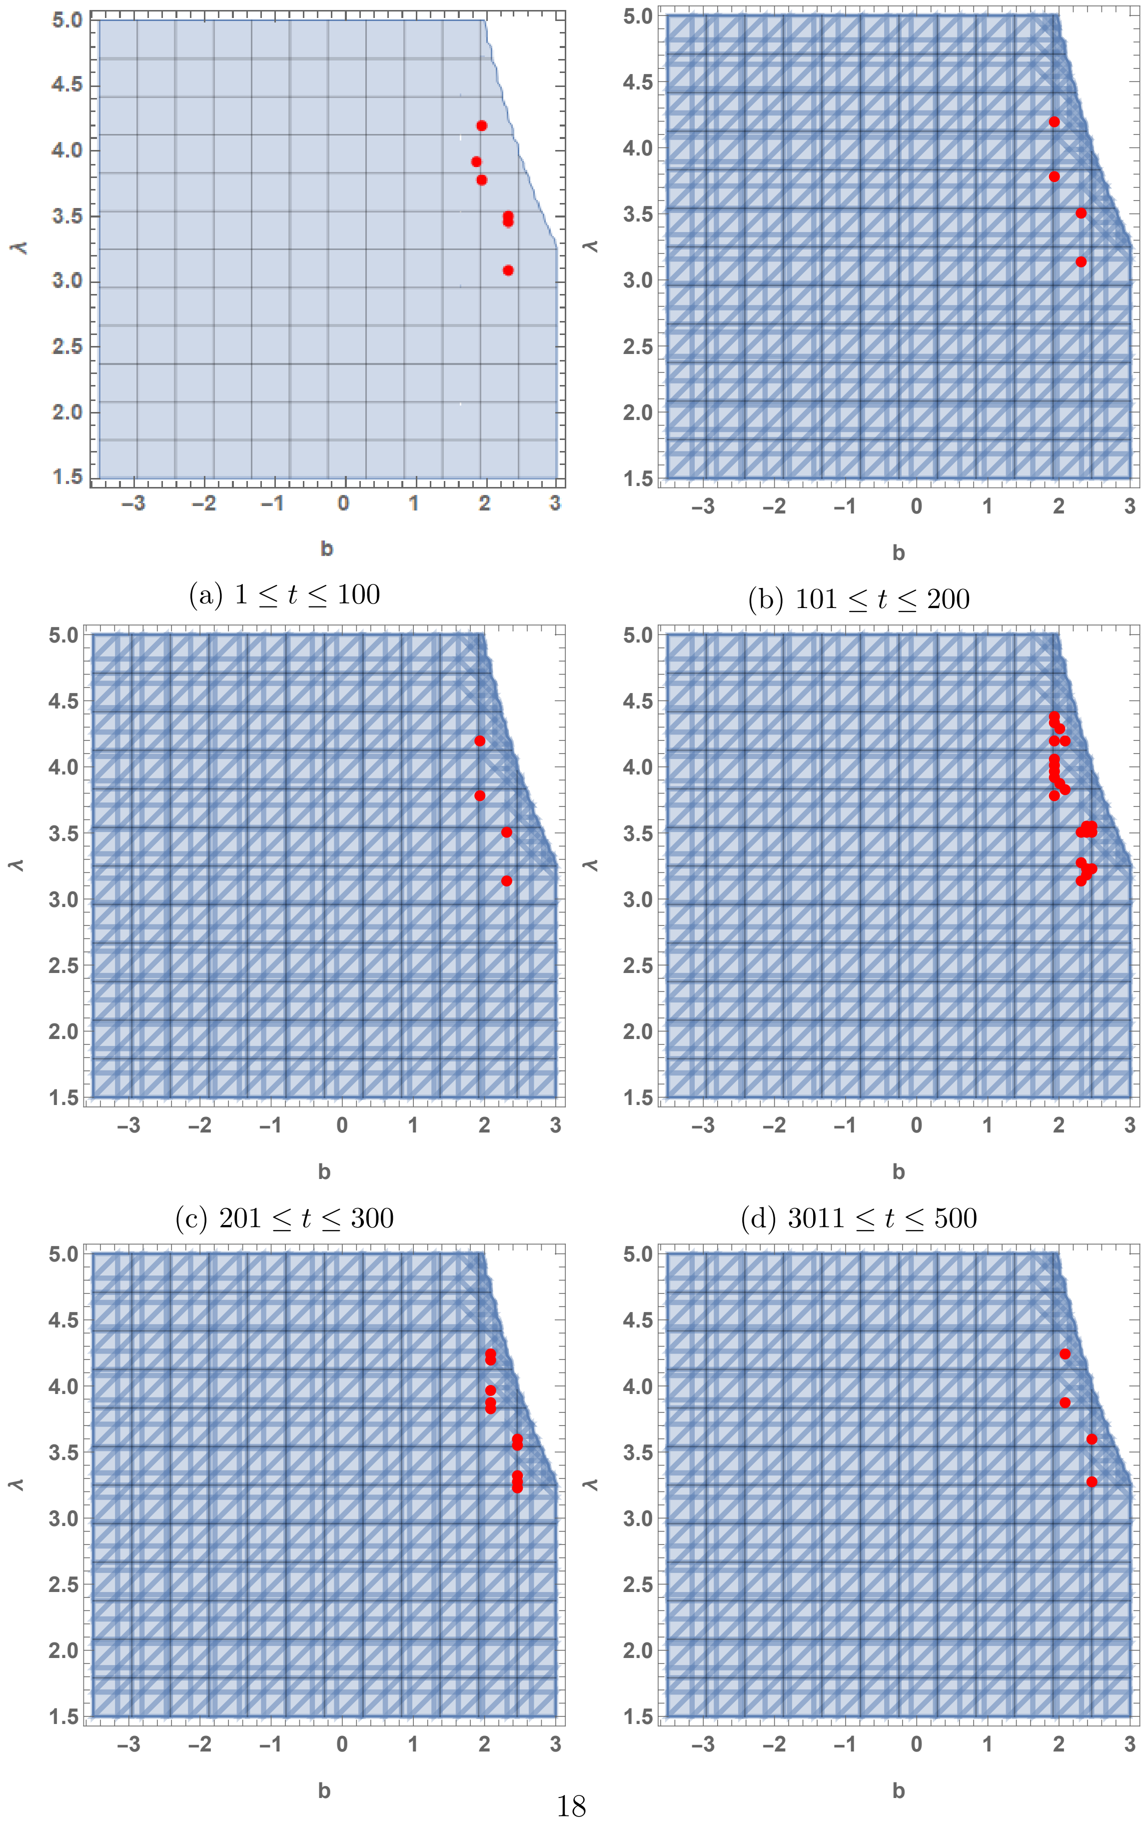

(a) $1 \leq t \leq 100$

(b) $101 \leq t \leq 200$

(c) $201 \leq t \leq 300$

(d) $3011 \leq t \leq 500$

(e) $501 \leq t \leq 1000$

(f) $1001 \leq t \leq 5000$

Figure 4: All points in the state space that are visited.

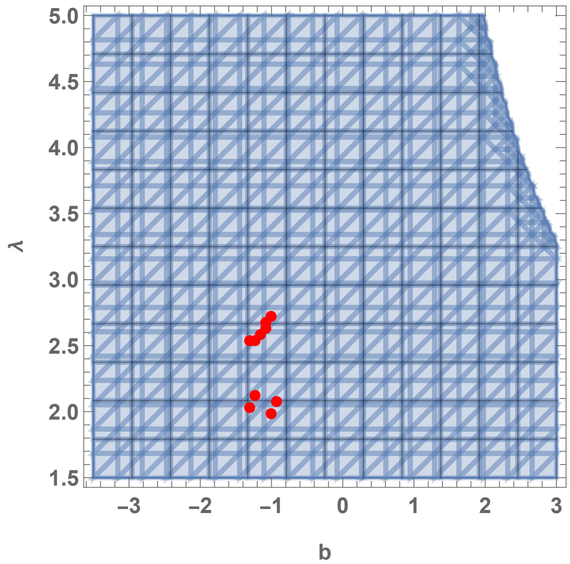

(a) $1 \leq t \leq 100$

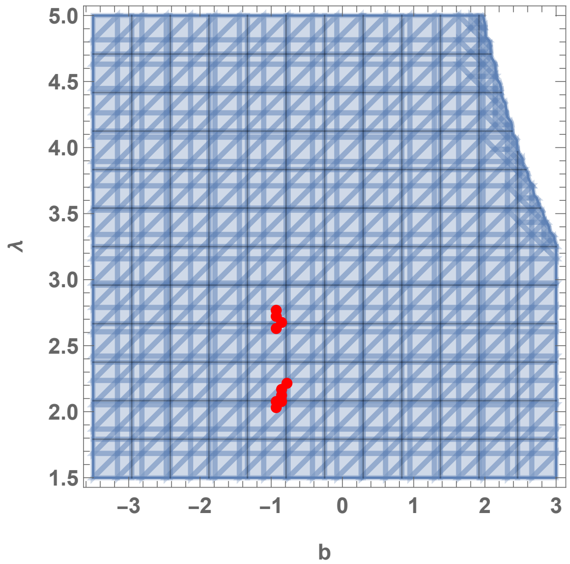

(b) $101 \leq t \leq 200$

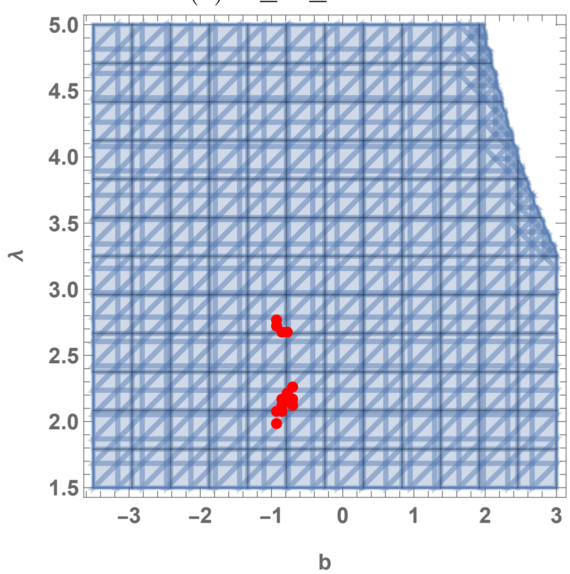

(c) $201 \leq t \leq 300$

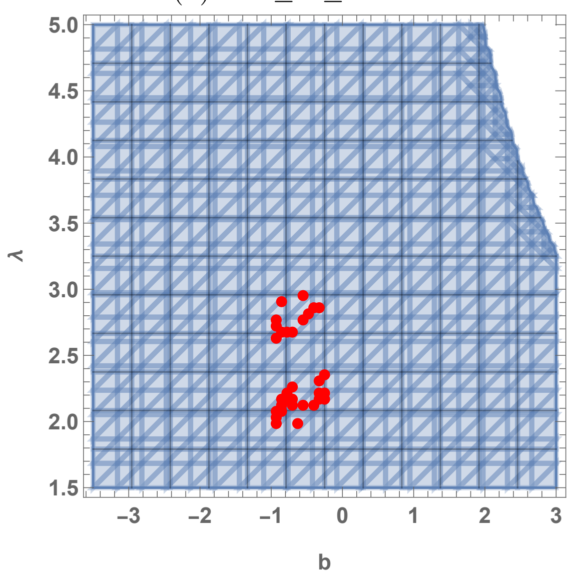

(d) $3011 \leq t \leq 500$

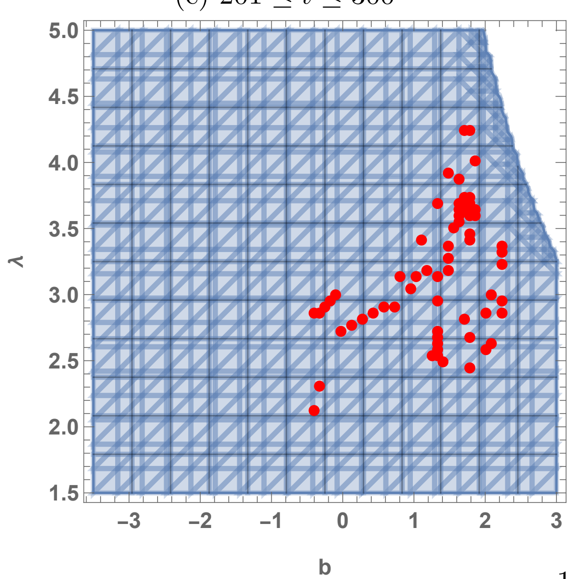

(e) $501 \leq t \leq 1000$

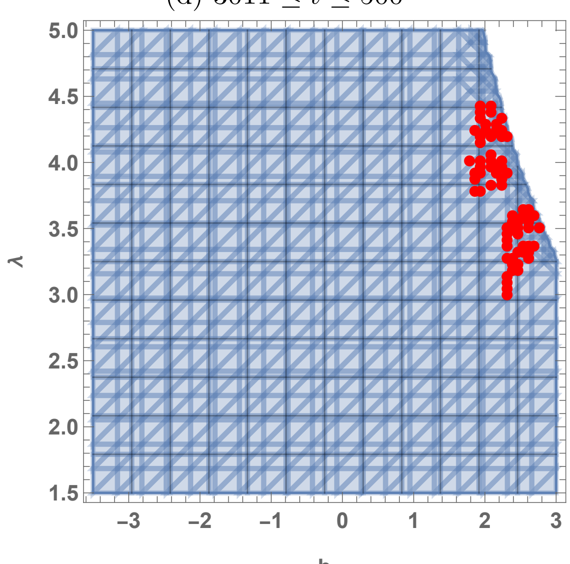

(f) $1001 \leq t \leq 5000$

Figure 5: All points in the state space that are visited.

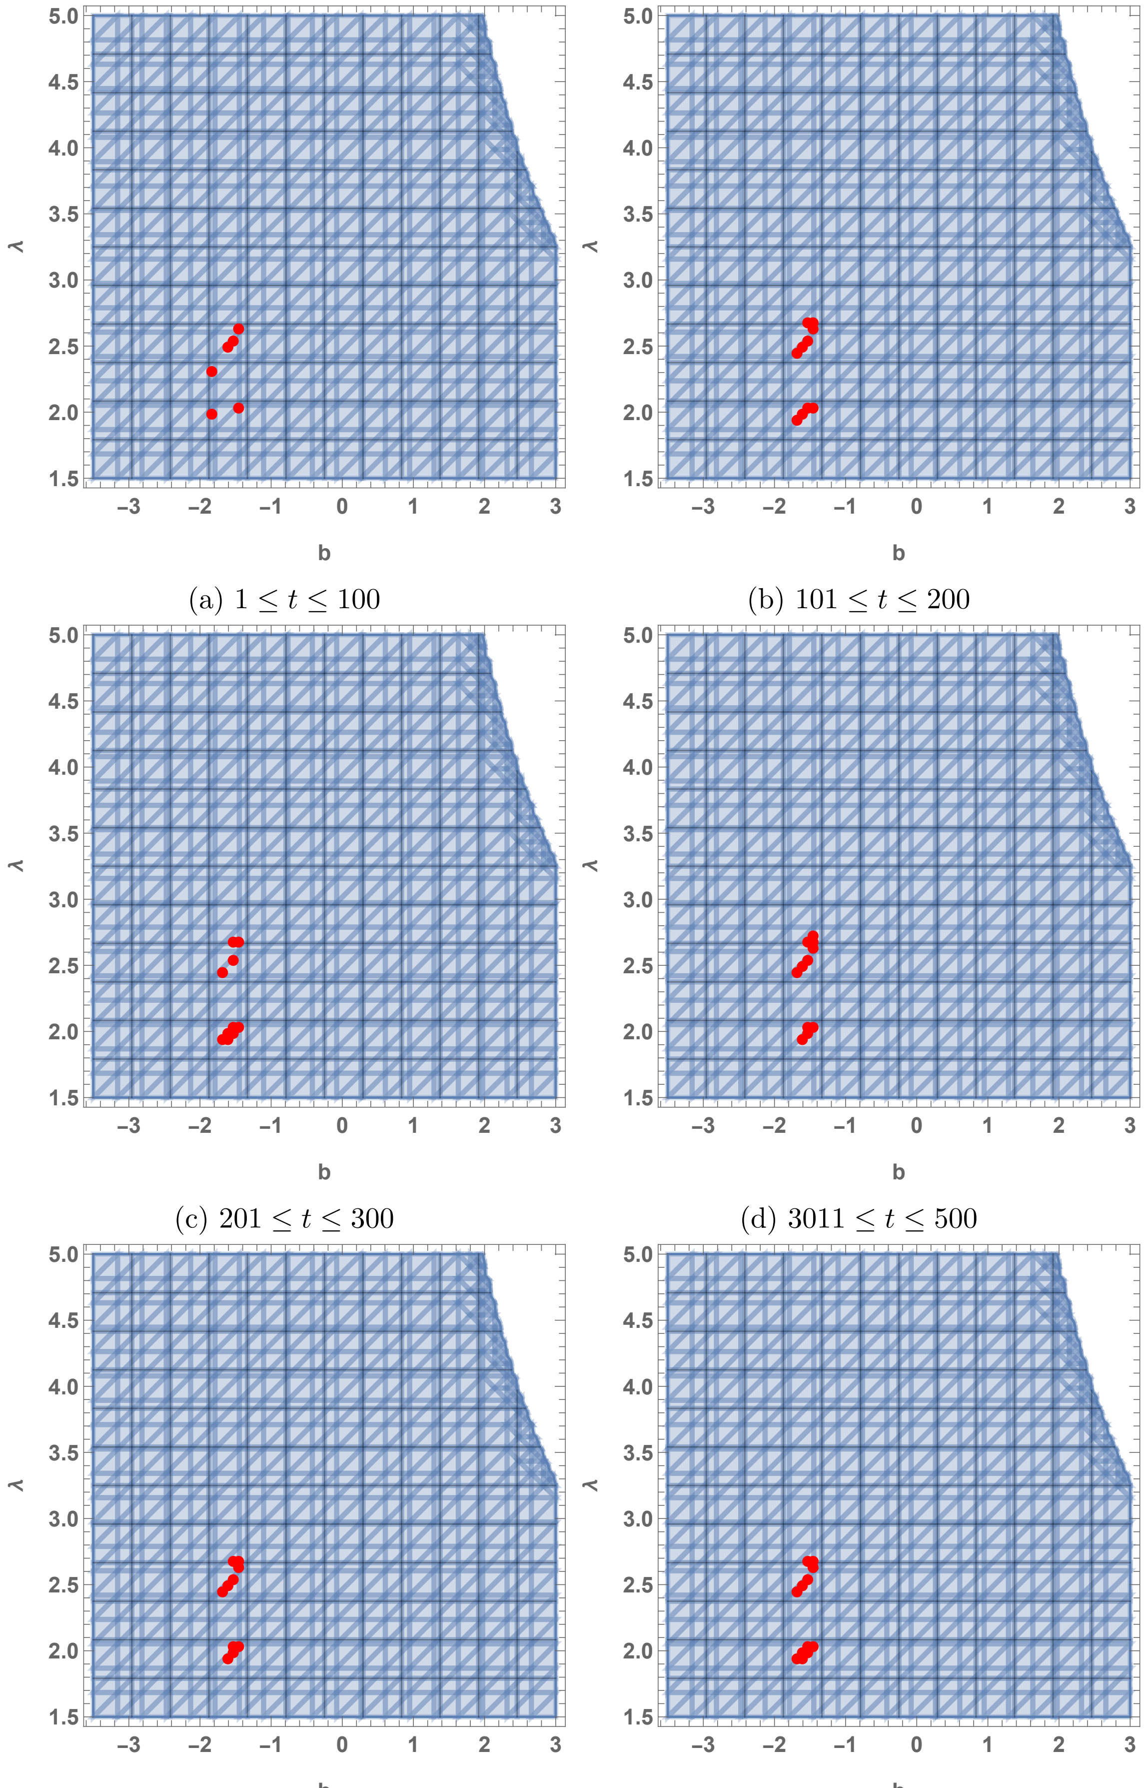

(a) $1 \leq t \leq 100$

(b) $101 \leq t \leq 200$

(c) $201 \leq t \leq 300$

(d) $3011 \leq t \leq 500$

(e) $501 \leq t \leq 1000$

(f) $1001 \leq t \leq 5000$

Figure 6: All points in the state space that are visited.

References

Aiyagari, S. , Marcet, A., Sargent, T., & Seppälä, J. (2002). Optimal Taxation without State-Contingent Debt. *Journal of Political Economy*, *110*(6), 1220–1254. doi: 10.1086/343744

Andersson, J. A. E., Gillis, J., Horn, G., Rawlings, J. B., & Diehl, M. (2018). CasADi – A software framework for nonlinear optimization and optimal control. *forthcoming in: Mathematical Programming Computation.*

Barro, R. J. (1979). On the Determination of the Public Debt. *Journal of Political Economy*, *87*(5, Part 1), 940–971. doi: 10.1086/260807

Buera, F., & Nicolini, J. P. (2004). Optimal maturity of government debt without state contingent bonds. *Journal of Monetary Economics*, *51*(3), 531–554. doi: 10.1016/j.jmoneco.2003.06.002

Judd, K. L. (1999). *Numerical Methods in Economics.* The MIT Press. doi: 10.1111/j.1467-6435.1999.tb02815.x

Kydland, F. E., & Prescott, E. C. (1977). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, *85*(3), 473–491. doi: 10.1086/260580

Kydland, F. E., & Prescott, E. C. (1980). Dynamic optimal taxation, rational expectations and optimal control. *Journal of Economic Dynamics and Control*, *2*, 79–91. doi: 10.1016/0165-1889(80)90052-4

Lam, S. K., Pitrou, A., & Seibert, S. (2015). Numba: a LLVM-based Python JIT compiler. , 7. doi: 10.1145/2833157.2833162

Lucas, R. E., & Stokey, N. L. (1983). Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics*, *12*(1), 55–93. doi: 10.1016/0304-3932(83)90049-1

Wächter, A., & Biegler, L. T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, *106*(1), 25–57. doi: 10.1007/s10107-004-0559-y

A Appendix

A.1 Exogenous Government Spending

In this section, we present additional results for the exogenous government spending case—the standard AMSS economy. Specifically, we look at the cases presented in Table 2.

Case #	σ_1	σ_2	$\underline{c}$	exogenous g	$\bar{g}$	b bounds	c bounds
5	1.0	1.0	0	No	$\{0.09, 0.27\}$	$[-10, 12]$	$[0.01, 1]$
6	1.0	1.0	0	No	$\{0.045, 0.135\}$	$[-5, 7]$	[0.01, 1]
7	1.0	1.0	0	No	$\{0.0225, 0.0675\}$	[]	[0.01, 1]

Table 2: Experiment setup

We study the feasible region for cases 5-7 with Π being a symmetric Markov transition matrix with a 50%, 66.67% and 80% probability of "staying-where-you-are", denoted as cases a, b, and c, respectively. I.e., case 5b denotes case 5 with the symmetric matrix with 66.67% "staying-where-you-are" probability. Figure 7 depicts the feasible regions.

Cases 5(a-c) and 6(a-c) are consistent with Section 5: The feasible regions are restricted to the second quadrant, i.e., the government cannot take up any debt. Changes in the Markov transition matrix do not qualitatively change this result. For very low government spending—in peace as well as in war—as in case 7(a-c), the results change qualitatively. It is now feasible for the government to borrow from the household and build up some debt. The same pattern holds as in the endogenous g case: with λ increasing, higher debt levels become infeasible.

[We are going to expand this by time-series simulations]

A.2 Endogenous Government Spending: Alternative Utility Functions

For alternative utility functions, we have not solved for closed-form solutions. In this case, we cannot employ GPUs efficiently and solve the dynamic programming problem on a cluster as alternative solution approach. In the following, we present the tools we use for the computation, the cases, and their implied feasible regions. In later versions of this paper, we analyse the time-series simulations for these cases as well.

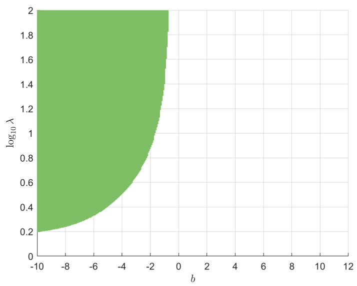

(a) Case 5a

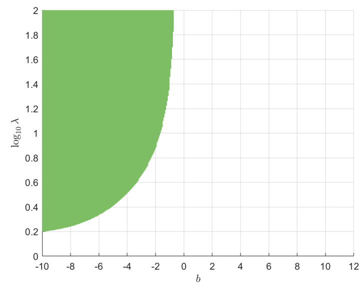

(b) Case 5b

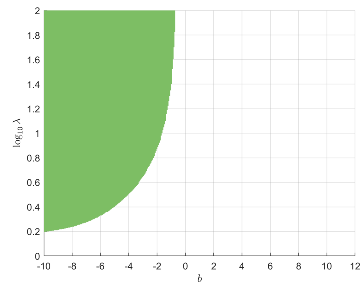

(c) Case 5c

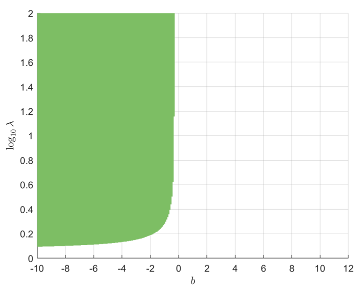

(d) Case 6a

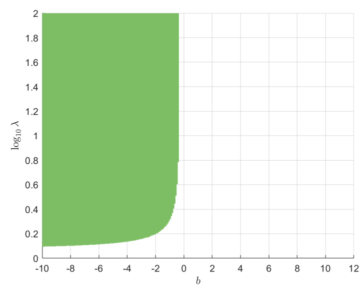

(e) Case 6b

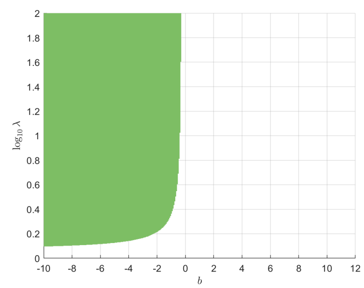

(f) Case 6c

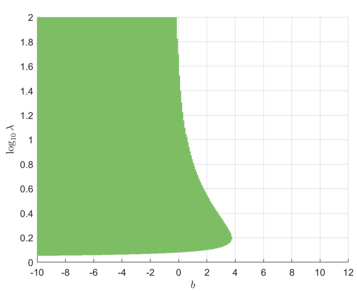

(g) Case 7a

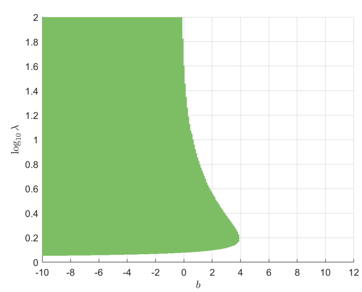

(h) Case 7b

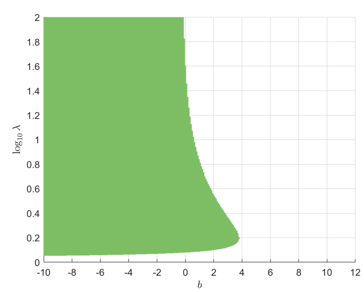

(i) Case 7c

Figure 7: Numerically feasible regions $\tilde{\Omega}(z)$ for cases in Table 2 and transition proability matricies Π for cases a, b, and c.

We have not solved for a closed-form solution for the case where either $\sigma_1 \neq 1$ or $\sigma_2 \neq 1$. Instead, we employ a combination of tools to efficiently solve the individual optimization problems. As optimizer, we use IPOPT by Wächter and Biegler (2006) in combination with the automatic differentiation tool CasADi by Andersson, Gillis, Horn, Rawlings, and Diehl (2018). The code is in written in Python, and we use Numba by Lam, Pitrou, and Seibert (2015) for just-in-time compilation. We rely on heuristics to accelerate the algorithm. In the future, we will approximate the solution to the individual optimization problem and use GPU computing for these utility functions as well.

We investigate the cases presented in Table 3. Recall the utility function from Equation

2

$$u(c,\ell,g,z) = \frac{(c+\underline{c})^{1-\sigma_1}}{1-\sigma_1} + \eta\frac{(1-\ell)^{1-\sigma_2}}{1-\sigma_2} - \theta(g-\bar{g}(z))^{\sigma_3}.$$

Case #	σ_1	σ_2	$\underline{c}$	exogenous g	$\bar{g}$	b bounds	c bounds
8	0.5	0.5	0.1	No	$\{0.09, 0.27\}$	$[-10, 12]$	$[0.01, 1]$
9	0.5	1.0	0.1	No	$\{0.09, 0.27\}$	$[-10, 12]$	$[0.01, 1]$
10	1.0	0.5	0.1	No	$\{0.09, 0.27\}$	$[-10, 12]$	$[0.01, 1]$

Table 3: Experiment setup

he utility functions are now combinations of sqrt utility functions and log utility functions. Figure 8 depicts the feasible regions in this case. The general feasible debt level is in case 8 and 9 lower than before, but qualitatively, they show the same pattern as before: The level of feasible debt shrinks with increasing λ.

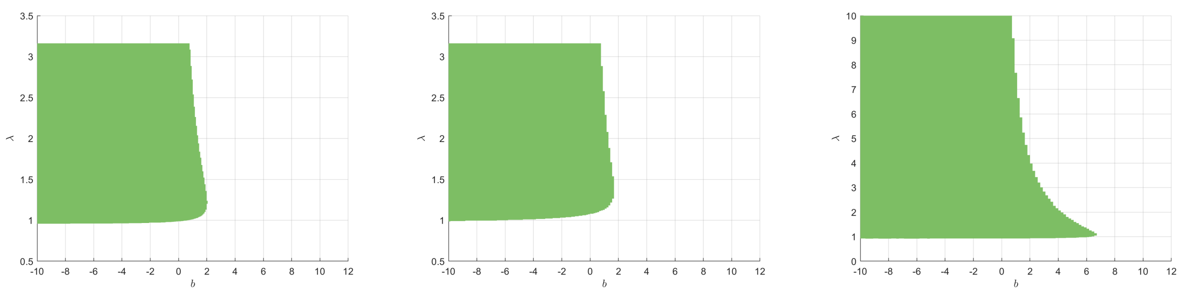

(a) Case 8 (b) Case 9 (c) Case 10

Figure 8: Numerically feasible regions $\tilde{\Omega}(z)$ for cases in 3