Optimal Dynamic Stochastic Fiscal Policy with Endogenous Debt Limits

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

April 13, 2022

Abstract

THIS IS A VERY ROUGH DRAFT. DO NOT QUOTE.

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 Barro (1979) and Aiyagari, Marcet, Sargent, and Seppälä (2002) incomplete market, dynamic fiscal policy model but introduce important changes. First, we assume that government spending is endogenous and implied by a utility function. Second, we use a dynamic programming formulation for the government's dynamic problem. Third, we use global optimization methods which can handle possible non-convexities. Fourth, we compute the endogenous debt limit without imposing any artificial debt limit. These new features lead to substantially different results compared to the previous literature. First, there is no tendency to accumulate a war chest large enough to allow taxation to disappear, even during long periods of peace. 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 can adjust its spending when faced with financial distress, it can sustain a credible reputation for honoring high debt levels. Fourth, as Judd has shown in a companion note, the Barro analysis assumes that the inefficiency cost of taxation is a quadratic function, a specification inconsistent with the fact that tax revenues are bounded. Microeconomically sound specifications of that function imply dramatically different long-run behavior. In fact, the robust result in the Barro model is that the optimal tax policy has the government accumulate assets to the level where taxes are minimal. The long-run policy

*Hoover Institution & NBER, kennethjudd@mac.com

†Department of Business Administration, University of Zurich, philipp.mueller@business.uzh.ch. Philipp Müller gratefully acknowledges the support of Kenneth Judd, Senior Fellow at the Hoover Institution.

‡Simon Business School, University of Rochester, syelteki@simon.rochester.edu

is not a random walk but instead a stable process with a target level of debt. Our analysis is fully consistent with tax theory and is also robust to changes in structural assumptions.

1 Introduction

The SARS-CoV-2 pandemic has disrupted the global economy. Public debt levels and deficits have reached record heights due to a contraction in real GDP and increased government expenses for economic relief and stimulus. In light of this alarming development, the ability of a fiscal system to finance government expenditures has become a major concern in public debates.

Central to these debates is the notion of a natural debt limit, i.e., the maximum level of public debt that can be sustained in the long run by fiscal policy and how tax policy should respond to increases in expenditures. Barro (1979) presented the original paper describing the time series process of optimal taxation. Under the assumption of a quadratic loss function, he showed that optimal tax policy would 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. He argues that "deficits are varied in order to maintain expected constancy in tax rates." [1]

However, the Barro analysis is flawed. A quadratic specification of the social cost of taxation allows unbounded levels of taxation and debt. Standard microeconomics implies that there is a maximum level of revenue and that the cost of taxation is infinite for levels beyond some finite bound. When his implicit assumption of a quadratic function is replaced with a fully specified micro model of taxation, the results are very different. In fact the robust result appears to be that the optimal tax policy drives the level of taxation to a low level, and asymptotically is a stable process instead of a random walk. This is exposited in a companion note by Judd.

Aiyagari et al. (2002)'s (AMSS) sought to put the Barro analysis on a microeconomic foundation incorporating Barro's assumptions of random expenditure shocks and incomplete markets. Economic agents in their model make fully rational decisions regarding consumption, labor supply, and asset accumulation. They presented examples where taxes roughly follow a random walk, as Barro argued. They also presented examples where high taxes in

[1] 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. This paper keeps the incomplete market structure in Barro, assuming that the only asset is the risk-free bond.

the short-run allow the government to accumulate assets so that in the long run, the tax rate is zero and all spending is financed out of its asset income (i.e., the government holds debt of the people). However, their approach had some weaknesses. First, it did not address the question of an endogenous, natural debt limit. In fact, its computational method never checked the feasibility of debt levels. Second, as is common in this literature, it assumed exogenous government spending. Third, it was based on only the first-order conditions of the government's dynamic optimization problem, ignoring the naturally occurring nonconvexities in optimal taxation models. Fourth, it assumed that government policy could be well approximated with a low-degree polynomial following the parameterized expectations algorithm (PEA) by Marcet (1988).

We extend the AMSS model by incorporating an endogenous government spending choice for the government and a taste shock Markov process to the preferred level of government spending. Note that this specification nests the original AMSS exogenous government spending specification in the case of penalizing any deviation of the preferred level by $-\inf$. We solve this model by using various computational methods that address the computational issues. Our computational approach allows for arbitrarily complex decision rules for the government and uses global optimization methods. The combination of more flexible computational tools and flexible government spending results in very different economic implications for the set of feasible debt levels and the dynamics of fiscal policy.

Our computational approach involves recasting the policy problem as an infinite horizon dynamic programming problem as first proposed by Kydland and Prescott (1980). They recognize the importance of determining the set of feasible policies and highlight that the state space over which the problem is defined is endogenous and must be determined jointly with solving the dynamic programming problem. They give an excellent discussion of the problems of finding the feasible region but also say "computing an optimal policy would appear quite formidable even for relatively simple parametric structures."We provide a set of computational techniques that can solve a wide variety of the kind of problems described in Kydland and Prescott (1980).

The presence of an endogenous feasible space presents multiple challenges. The government's value function will have a very high curvature and diverge as debt approaches its endogenous limit. The resulting dynamic programming problem presents several computational challenges and necessarily makes heavy demands on computational resources, which explains why this approach is generally not taken. In particular, AMSS combines the PEA method with the first-order approach to dynamic principal-agent problems. [2] Using our combination of computational tools and more general economic assumptions, we re-address

[2] A companion paper will describe in detail several problems with the AMSS solution method.

questions regarding optimal taxation and debt management in a more realistic and flexible framework. 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 under the more general framework of endogenous government spending and endogenous debt limits have substantially different implications than earlier analyses. First, the behavior of an optimal policy is, over long horizons (e.g., 1000 years), much more complex than simpler models imply. The Barro random walk result implies, as noted by Bizer and Durlauf (1990), that either the martingale is degenerate or somebody's budget constraint is asymptotically violated. In our analysis, the economic fundamentals determine limits on debt without violating any budget constraints. Also, the stochastic process describing debt and taxes is not captured well by any simple auto-regressive process. In particular, the long-run distribution of debt is multimodal, and the long-run level of debt is history-dependent. If the 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, 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 expenditure remains low forever. When debt is high, governments will avoid default by reducing spending and using taxes to finance a persistently high debt, but not enough to reduce the debt.

When we use our computational methods to examine the AMSS case of exogenous, fixed government spending, we find results different from the AMSS model and their computational methods. In particular, we find that if spending shocks are similar to US experience, then no positive level of debt is feasible. The key fact is that if spending is exogenous, then feasibility means that there is a sequence of tax rates that will fund both the spending and interest burden for any sequence of spending shocks. In the case of US-like experience, if government debt is initially positive, but the government is hit with a sequence of high spending shocks that it cannot control, such as a 30-year version of WWII, then the government cannot finance its obligations unless it already had a war chest. Assuming fixed exogenous government spending will imply that any level of debt is infeasible for many developed countries.

2 US Fiscal Policy in the past 120 years

The US, and many other developed countries, carry significant levels of debt. The next figures display the history of public debt, federal government expenditures, and receipts since 1900 for the US. Figure 1 shows the war-related peaks, a small one for WWI and a very large one for WWII. Note that there are no peaks for either the Korean or Vietnam wars. Figure 2 shows the time series for expenditure and receipts. We see a permanent increase in revenue after the peak in WWII spending, as predicted by tax smoothing arguments. Neither the Korean nor Vietnam war generates a major spike in spending. In fact, government spending increases after the end of the Vietnam War. It is only in the mid-1990's that spending falls relative to revenues.

The major shocks in the past 20 years have been the Great Recession and the recent COVID-19 pandemic, which together have pushed the level of debt to levels not seen since the end of WWII. The fiscal actions since mid-2020 will substantially increase spending and push debt levels to an unprecedented level. These developments cause major concerns about the sustainability of the US debt.

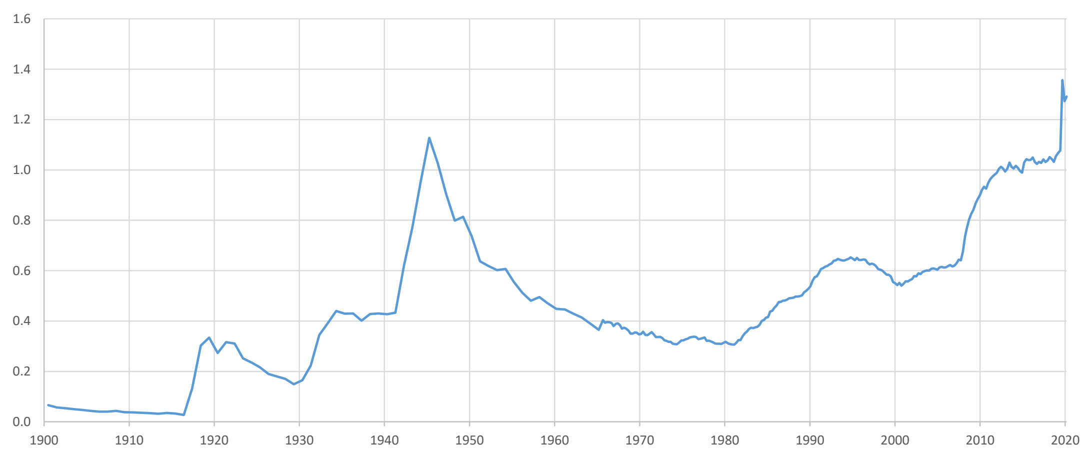

Figure 1: Total public debt as percent of gross domestic product.

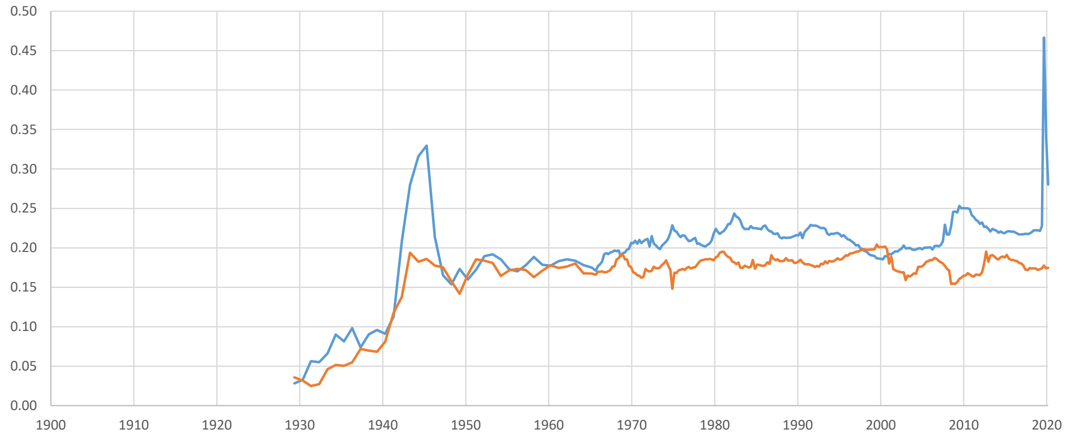

Figure 2: Federal expenditures (blue) and receipts (orange) as percent of gross domestic product.

We depart from much of the macroeconomic literature by assuming that spending is flexible. We argue that any examination of US history contradicts the assumption of exogenous spending even when it comes to major wars. WWI was clearly a voluntary war. Germany and its allies presented no threat to the US or the British Empire. In "The Pity of War", Niall Ferguson argues that Great Britain was not compelled by the 1839 Treaty of London to intervene when German armies moved through Belgium to attack France. Acting as a private citizen, Herbert Hoover, through his Commission for Relief in Belgium, convinced Germany of Belgium's special status and allowed the shipment of food to Belgium, showing that Germany did not treat Belgium as an enemy. In 1917, many, most notably Senator Robert M. La Follette of Wisconsin, opposed American participation, arguing that the US should maintain its traditional aversion to involvement in European wars. Perhaps entering WWI was the correct decision but it was a choice.

American participation in WWII was arguably at least partly forced by events. The US would surely have responded to both the attacks on Pearl Harbor and the Philippines as well as ongoing attacks by German submarines on US naval vessels.[3] It was natural for the US to come to the aid of Great Britain and form an Anglo-American alliance to defend British and American interests against German expansion. However, even before Dec. 7, 1941, the US fully embraced Churchill's objective to end German domination of Europe, a choice predicated on the belief that the US would join the UK in that goal.[4] The magnitude

[3] In October, 1941 the USS Kearney dropped depth charges on German U-boats as part of an effort to defend a British convoy. On October 17, U-568 hit the Kearney with a torpedo, killing 11 sailors. On October 31, 1941, the US Navy destroyer USS Reuben James was sunk by U-552, with the loss of 100 sailors.

[4] In March, 1941, FDR approved ABC-1, the report of a secret meeting involving American, British, and

of WWII spending was chosen, not forced by necessity.

After WWII, war spending does not produce any major peaks. Defense spending was persistently high for a few decades but has dropped to only roughly two percent of GDP. Movements in debt and spending reflected decisions in various categories of flexible domestic spending. Considerations of fiscal and military history of the US illustrate clearly that any analysis of fiscal policy that wants to examine historical fiscal policy should recognize that government spending is chosen, not forced by exogenous circumstances.

Moreover, the assumption of exogenous spending is not an acceptable simplification. It is clear from our computations that it would be very difficult to specify an exogenous spending process that calibrates the US experience, put it into the AMSS model and find positive, persistent feasible debt levels. However, if we assume flexibility in spending, persistent moderate debt levels can be observed. The difference is that if hit with an unlikely persistent sequence of positive spending shocks, our model will allow the government to reduce spending as it approaches infeasible debt levels. In contrast, the AMSS model could not avoid default. We illustrate this in examples where spending is only slightly flexible. The conclusion is rather obvious: the feasibility of debt depends critically on the ability to remain solvent even when faced with a very low probability sequence of adverse events.

The key assumption in our paper, as in Barro and AMSS, is that governments cannot default on debt. The US Constitution requires the US government to pay its debts, and the unwritten British constitution appears to do the same for the UK. Inflation has been used by both the US and the UK to partially default on debt, but that is best modeled by incorporating both taxation and inflation and their distinct distortionary costs in a more complex model. We assume that government debt is risk-free because it allows us to compare our results to other work, and there appears to be a consensus that this is a reasonable first-order assumption to make for the US, UK, and some other major economies.

3 A Model of Optimal Taxation with only Risk-free Debt

The model is a simple one. The economy is inhabited by a government and a continuum of identical infinitely-lived households. Households are endowed with one unit of time in each period, and provide labor, ℓ, to produce market consumption goods, c, or public goods, g. The supply of ℓ_t is limited to the closed interval $[0, 1]$ and time not spent in formal labor activities, that is, $1 - \ell$, is spent at home dedicated to leisure activities or home production.

Canadian military staff that outlined the plan for the European theatre. ABC-1 reads like a history of WWII in Europe. The Lend-Lease act was also passed, which provided Stalin with enormous amounts of military aid between the Summer of 1941 and May, 1945, another expensive choice made by the US.

The government imposes a proportional tax on labor income, issues a transfer to households, sells risk-free bonds held by the households, and spends g on public goods. Households spend their income on consumption c and bonds b which are held for one period.

We provide the details of our model in the following and highlight the differences to the AMSS model.

3.1 Households

Utility for a representative household is a function of consumption, c, labor supply, ℓ, and government expenditures, g. We denote the utility by $u(c, \ell, g, z)$ and assume that it is increasing in consumption, $\frac{\partial u}{\partial c} > 0$, and decreasing in time spent working, $\frac{\partial u}{\partial \ell} < 0$.

The taste for government spending depends on taste shocks, z, which follow a finite Markov process with a transition matrix $\pi(z'|z)$. We assume that the utility u is a smooth function of g penalizing deviations from the preferred level of government spending, $\bar{g}(z)$. In contrast to our model specification, the government spending in AMSS is exogenous and fixed to the preferred level of government spending. This would translate in our framework to penalizing any deviation of g from $\bar{g}(z)$ by ∞ unless g equals $\bar{g}(z)$. We refer to the AMSS specification as the fixed-g case and our specification as the flexible-g case.

Households value the consumption, labor, and government expenditure stream, $\{c_t, \ell_t, g_t\}_{t=0}^{\infty}$, by its present value,

$$E_0 \sum_{t=0}^{\infty} \beta^t \, u(c_t, \ell_t, g_t, z_t), \tag{1}$$

where the discount factor β is in the open interval $(0, 1)$ and E_0 denotes the expectation conditioned on information at time 0.

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 $p_{b,t}$. In each period t, b_t is the number of bonds maturing at the beginning of the current period, and b_{t+1} the number of bonds issued today and maturing at period $t + 1$. When $b_t > 0$, the government is in debt to the households, when $b_t < 0$, the households are in debt to the government. Households pay a time-varying flat rate tax, τ_t, on their labor income and receive a transfer payment, $tr_t \geq 0$. The budget constraint of a household in period t,

$$(c_t + p_{b,t} b_{t+1}) - (b_t + tr_t + (1 - \tau_t)\ell_t) \leq 0. \tag{2}$$

The households accept the fiscal policy, $\Phi_t \equiv \{\tau_t, tr_t, g_t\}$, and taste shocks z_t as exogenous variables, while their bond holding choice, b_{t+1}, is endogenous. Households maximize their expected discounted payoff,

$$\max_{\{b_{t+1},c_t,\ell_t\}_{t=0}^{\infty}} E_0 \left\{ \sum_{t=0}^{\infty} \beta^t u(c_t, \ell_t, g_t, z_t) | \Phi_t \right\}, \tag{3}$$

given fiscal policy Φ and taste shock z, subject to the sequence of intertemporal budget constraints,

$$\begin{aligned} (c_t + p_{b,t} b_{t+1}) - (b_t + tr_t + (1 - \tau_t)\ell_t) \leq 0 \qquad & \forall t \\ c_t, \ell_t \geq 0 \qquad & \forall t. \end{aligned}$$

We derive the Lagrangian,

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

where λ_t denotes the Lagrange multiplier on the t-period budget constraint. For expositional purposes, we assume that the non-negativity conditions do not bind in the following. The household's first-order conditions,

$$\frac{\partial}{\partial c} : -\lambda_t + \frac{\partial}{\partial c} u(c_t, \ell_t, g_t, z_t) = 0 \tag{4}$$

$$\frac{\partial}{\partial \ell} : (1 - \tau_t)\lambda_t + \frac{\partial}{\partial \ell} u(c_t, \ell_t, g_t, z_t) = 0 \tag{5}$$

$$\frac{\partial}{\partial \lambda_t} : -b_t + c_t - tr_t - \ell_t(1 - \tau_t) + p_{b,t} b_{t+1} = 0 \tag{6}$$

$$\frac{\partial}{\partial b_{t+1}} : \beta \sum_{z_{t+1}} \lambda_{t+1}(z_{t+1}|z_t)\pi(z_{t+1}|z_t) - p_{b,t}\lambda_t = 0. \tag{7}$$

The Euler equation (7) is a function of both the current shadow price λ_t and the next period shadow price $\lambda_{t+1}(z_{t+1}|z_t)$. Hence, the household's labor supply, consumption, and debt decisions, ℓ_t, c_t and b_{t+1}, respectively, depend also upon future fiscal policy and not only upon the current state of the economy, z_t, and current fiscal policy Φ_t. When the government determines its policies, it must take into consideration the effect of its future policies on households' behavior in earlier periods.

3.2 Government

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

$$(\tau_t l_t + p_{b,t} b_{t+1}) - (b_t + tr_t + g_t) = 0. \tag{8}$$

Recall that b_t denotes payouts of debt at the beginning of period t.

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

$$(1 - \ell_t) + c_t + g_t = 1. \tag{9}$$

The timing of the moves is as follows: After the realization of the current taste shock z_t, the government makes its tax, transfer, and spending decisions, τ_t, tr_t, and g_t, chooses market price for bonds $p_{b,t}$, the shadow price of consumption for the next period, and recommends allocations for the household c_t, ℓ_t, b_t. The households solve their own optimization problem, given the fiscal policy choice and pick the allocation suggested by the government if it is in their interest to do so. To ensure that households follow through with his plan, the government chooses fiscal policy and household allocations that are consistent with households' optimal choices of consumption and labor. Additionally, these policies must deliver the shadow price λ_t and debt b_t from the household's problem.

The government's problem can be written as:

$$\max_{c_t, \ell_t, g_t, \Phi_t} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t, \ell_t, g_t, z_t) \tag{10}$$

subject to its budget constraint

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

aggregate resource constraint

$$\ell_t - c_t - g_t = 0,$$

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

$$-\lambda_t + \frac{\partial}{\partial c} u(c_t, \ell_t, g_t, z_t) = 0 \tag{11}$$

$$(1-\tau_t)\lambda_t + \frac{\partial}{\partial \ell} u(c_t, \ell_t, g_t, z_t) = 0 \tag{12}$$

$$-b_t + c_t - tr_t - \ell_t(1-\tau_t) + p_{b,t} b_{t+1} = 0 \tag{13}$$

$$\beta \sum_{z_{t+1}} \lambda_{t+1}(z_{t+1}|z_t)\pi(z_{t+1}|z_t) - p_{b,t}\lambda_t = 0 \tag{14}$$

$$\tag{15}$$

non-negativity constraints,

$$c_t, \ell_t, g_t, p_{b,t}, \lambda_t, \lambda_{t+1}(z^+|z_t), tr_t \geq 0.$$

4 Problem Formulations

In this section, we first present the first-order approach taken by, for example, Aiyagari et al. (2002) and point out potential shortcomings of this approach. Next, we present Kydland and Prescott (1977) on why standard dynamic programming is inappropriate for economic planning problems and finally, a reformulation of the dynamic programming approach that applies to economic planning problems presented in Kydland and Prescott (1980), which we apply to our model.

4.1 First-order approach

The first-order approach starts from the government's optimization problem. We denote the utility $u(c_t, l_t, g_t, z_t)$ by U_t, the government budget equation (8) by GB, the production possibility frontier by PPF, and the consumer incentive compatibility constraints—that is, the first-order conditions of the household's problem with respect to consumption, labor, next period's bond holdings, and the shadow price of consumption—by IC^c, IC^ℓ, IC^b, and IC^λ, respectively. For notational brevity, U_{t,x_t} denotes the partial derivative of U_t with respect to x_t; analogously, $U_{t,x_t x_t}$ denotes the second order partial derivative of U_t with respect to x_t.

The Lagrangian of government problem is

$$L = E_0\Bigg[\sum_{t=0}^{\infty}\beta^t U_t - \theta_t^{GB} GB_t - \theta_t^{PPF} PPF_t - \\ \varphi_t^c IC_t^c - \varphi_t^\ell IC_t^\ell - \varphi_t^b IC_t^b - \varphi_t^\lambda IC_t^{\lambda_t}\Bigg]. \tag{16}$$

The first-order conditions with respect to the consumer's choices c_t, ℓ_t, and b_{t+1} and the government's choices

$$\begin{aligned}
\frac{\partial}{\partial c_t} &: \quad U_{t,c_t} - \theta_t^{PPF} - \varphi_t^c U_{t,c_t c_t} - \varphi_t^\ell U_{t,\ell_t c_t} - \varphi^\lambda = 0 \\
\frac{\partial}{\partial \ell_t} &: \quad U_{\ell_t} - \varphi_t^{GB}\tau_t - \varphi_t^{RC} - \varphi_t^c U_{t,c_t\ell_t} - \varphi_t^\ell U_{t,\ell_t\ell_t} - \varphi_t^\lambda(1-\tau_t) = 0 \\
: \frac{\partial}{\partial b_{t+1}} &: - p_{b,t}\theta^{GB} + \beta\theta^{GB} = 0 \\
\frac{\partial}{\partial g_t} &: \quad U_{t,g_t} + \varphi_t^{GB} + \varphi_t^{RC} - \varphi^c U_{t,c_t g_t} - \varphi^\ell U_{t,\ell_t g_t} = 0 \\
\frac{\partial}{\partial \tau_t} &: \quad -\varphi_t^{GB}\ell_t + \varphi_t^\ell \lambda_t = 0 \\
\frac{\partial}{\partial \lambda_t} &: \quad \varphi^c - \varphi^\ell(1-\tau_t) + \varphi^{EE} p_{b,t} = 0
\end{aligned}$$

household's willingness to substitute consumption over time not only depends on c but also on ℓ.

This is not a convex problem. This failure makes the first-order approach problematic because the first-order conditions may have multiple solutions, some of which are not even local optima. This necessitates global optimization methods.

4.2 Recursive formulation

Standard control theory is appropriate for systems whose state movement only depends on current an past policy decisions as well as the current state. In the optimal tax problem above, Constraint (14) violates this assumption. It depends on today's as well as tomorrow's shadow price of consumption. In this case, the optimal policy is time-inconsistent in the sense that an optimal policy sequence chosen at time 0, $\{\Phi_t^0\}_{t=0}^{\infty}$, is not optimal at time $s > 0$. However, as Kydland and Prescott (1977) note, standard control theory solves for the time-consistent sub-optimal solution. We drop the t subscript for notational brevity in the following and denote the next period by a + superscript.

An alternative formulation of the recursive problem by Kydland and Prescott (1980) allows the application of control theory for economic planning problems: This optimization problem can be reformulated into a recursive formulation by taking the current bond level, the marginal utility of consumption, and the current shock, that is, (b, λ, z), as state variables. These suffice as sufficient statistic for the dynamic programming problem. The government's dynamic programming problem reads:

$$\max_{c,l,b,g,\tau,\lambda,\lambda^+} \quad u(c, l, g, z) + \beta \mathbb{E}[V(b^+, \lambda^+(z^+|z), z^+)] \tag{17}$$

subject to its budget constraint

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

aggregate resource constraint

$$\ell = c + g,$$

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

$$-\lambda + \frac{\partial}{\partial c} u(c, \ell, g, z) = 0 \tag{18}$$

$$(1 - \tau)\lambda + \frac{\partial}{\partial \ell} u(c, \ell, g, z) = 0 \tag{19}$$

$$-b + c - tr - \ell(1 - \tau) + p_b b^+ = 0 \tag{20}$$

$$\beta \sum_{z^+} \lambda^+(z^+|z)\pi(z^+|z) - p_b\lambda = 0 \tag{21}$$

and the non-negativity constraints,

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

Not all state variables $\{b, \lambda\}$ allow for a policy sequence with equilibrium, that is, they do not allow for a policy sequence such that the household's incentive compatibility, the government budget constraints, and the aggregate resource constraint are fulfilled. We denote the set state variables allowing for such a policy sequence by $\Omega(z)$ and refer to it as feasible set. In general, there exists no closed-form solution for the feasibility set $\Omega(z)$; in Section 4.3, we describe a recursive algorithm to approximate it.

Solving the dynamic programming defined by Equation (17) must be done carefully which we address in the next section.

4.3 Computational Algorithm

Solving the stochastic dynamic policy problem of the government presents many computational challenges: (i) only a subset of the computational grid Γ is economically feasible, the set of feasible states $\Omega(z) \subseteq \Gamma$, (ii) $\Omega(z)$ is unknown and may be non-convex, which complicates the approximation of the value function because the value function is likely to diverge close to the boundary, (iii) the value function close to infeasible states will be highly curved and diverge to -∞ as we move to the boundary of the feasible set and (iv) there is no guarantee that the value function is concave. These issues present nontrivial computational challenges which we address in this section.

4.3.1 Feasible Region

We solve for the numerically feasible region $\tilde{\Omega}(g) \subseteq \Omega(g)$ by iteratively improving the k-th approximation of the feasible region $\tilde{\Omega}^k(g)$ until we find its fixed point, i.e., $\tilde{\Omega}^{k+1}(g) = \tilde{\Omega}^k(g)$.[5] The implementation of this process is straightforward: As a starting guess, we choose a "stay where you are" policy function, i.e., $b^+ = b$ and $\lambda^+ = \lambda$, and solve solve the government's optimization problem. This yields the initial feasible set $\tilde{\Omega}^0(g)$. We improve upon the initial guess $\tilde{\Omega}^0(g)$ by selecting for each infeasible state a state transition to a state that turned feasible in the previous iteration, producing $\tilde{\Omega}^1(g)$. This iteration continues until we have reached a fixed point, i.e, until $\tilde{\Omega}^k(g) = \tilde{\Omega}^{k+1}(g)$.[6]

4.3.2 Discretization

The feasible region turns out to be non-convex, and close to its boundary, the value function diverges. This poses great challenges to the functional approximation thereof. Thus, we apply discrete state dynamic programming and discretize the state variables, b and λ. We apply the same discretization to the value function and policy function.

We use a rectangular grid and uniformly discretize b. It may appear natural to discretize the promised marginal utility of consumption, λ, uniformly as well. However, for utility functions satisfying the Inada conditions, λ tends towards infinity for $c \to 0$. Thus, instead of discretizing λ uniformly, we discretize the consumption c uniformly starting from some lower bound c_{min} up to 1. This leads to a grid which is non-uniformly spaced in λ with

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

[6]We have also developed a complementary procedure that begins with some clearly infeasible states and then find other clearly infeasible states. Preliminary results indicate that the two procedures produce the same feasible set.

$\lambda \in [uc'(1), uc'(c_{min})]$. We use λ as the state variable but discretize it so that the discretization in consumption is uniform.

The state variables, b and λ, are (almost) continuous and thus require a fine discretization. In addition to the implementation presented in Section 4.3.3, we rely on an adaptive grid refinement. We begin with a coarse discretization which helps identify regions of the state space that are purely transient; that is, it is feasible to begin at those states but the solution says that the optimal policy will never revisit those states. With this information, we can then use a finer grid on the relevant portion of the state space.[7]

4.3.3 Algorithm and Implementation

Initially, we start with determining the feasible region as presented in Section 4.3.1. In short, we start with the initial policy function "stay where you are", and iteratively search for each infeasible state for a feasible policy until the feasible region converges.

Once we have determined the set of feasible states, we need to solve the dynamic programming problem. Note that the dynamic programming problem in Equation (17) can be seen as a bi-level optimization problem, in which we at first solve the inner, static problem having (b^+, λ^+) fixed:

$$
\begin{aligned}
U(b^+, \lambda^+) = \max_{c,l,p,tr,\tau,g} \quad & u(c, \ell, g, z) \\
s.t. \quad (\tau\ell + p_b b^*) - (b + tr + g) &= 0 \\
\ell - c - g &= 0 \\
-\lambda + \frac{\partial}{\partial c} u(c, \ell, g, z) &= 0 \\
(1-\tau)\lambda + \frac{\partial}{\partial \ell} u(c, \ell, g, z) &= 0 \\
-b + c - tr - \ell(1-\tau) + p_b b^+ &= 0 \\
\beta \sum_{z^+} \lambda^+(z^+|z)\pi(z^+|z) - p_b\lambda &= 0 \\
c, \ell, g, p_b, \lambda^*, \lambda^+(z^+|z), \mu, tr &\geq 0
\end{aligned}
$$

which allows for a (approximate) closed-form solution of the inner, static problem. The outer,

[7] Preliminary computations indicate that we have reached a discretization fine enough that an interpolant of the discrete data satisfies the optimality conditions for the value function over the continuous space. Ultimately, we will have enough data to determine good functional forms for the continuous state problem.

dynamic problem reads

$$V(b,\lambda,g) = \max_{b^+,\lambda^+} \; U(b^+,\lambda^+) + \beta\mathbb{E}[V(b^+,\lambda^+,g^+)] \qquad (22)$$
$$s.t. \;\; (b^+,\lambda^+) \in \Omega(g^+),$$

where $U(b^+,\lambda^+)$ denotes the maximized utility from the inner, static problem for fixed b^+ and λ^+. We rely on this formulation in the following.

To solve the dynamic programming problem for a fixed state (b,λ,g), we apply a policy improvement step. In this policy improvement step, we globally search for the $(b^+,\lambda^+) \in \Omega(g^+)$ that maximizes (22). The closed-form solution of the inner, static problem allows for an implementation of the global search that employs graphical processing units (GPUs). We apply this global search for all $(b,\lambda,g) \in \Omega(g)$. The policy iteration continues until the L^1 differences between successive value functions are less than 10^{-13}.

An important issue is whether our discretization is sufficiently fine for the results to be close to the results of the true problem where the states are continuous. We tried alternative discretizations and found that the 500x500 discretization which is used in our results below was not significantly different from coarser discretizations in terms of the fundamental economic insights. It is also clear that all of our results are consistent with agents making ϵ-optimal for small values of ϵ. The results are at least consistent with agent optimization at a level of accuracy better than what most believe to be the quality of actual optimization in the real world.

5 Results

In this section, we solve our model with endogenous government spending and contrast its results to the results implied by the Aiyagari et al. (2002) model, where the government spending is exogenous. Central to our comparison is the endogenous feasibility region $\Omega(z)$, which comprises the set of (b,λ) points that allow for a policy sequence with equilibrium.

5.1 Parameterization

We will follow the literature by assuming an additively separable utility in consumption utility, u_c, labor utility, u_l, and government expenditure utility, u_g, taking the form:

$$u(c,\ell,g,z) = \frac{(c+\underline{c})^{1-\sigma_1}}{1-\sigma_1} + \eta\frac{(1-\ell)^{1-\sigma_2}}{1-\sigma_2} - \theta(g-z)^{\sigma_3}, \qquad (23)$$

where θ, η, and $\underline{c} \in \mathbb{R}^+$. A $\underline{c} > 0$ allows (market) consumption, c, to equal 0 in the presence of high taxes due to a finite marginal utility of consumption for $c = 0$. This choice, in turn, implies that the revenue-maximizing tax rate is strictly less than unity.

The additively separable utility function specification (23) allows for a closed-form solution of the household problem. Equations (24-26) neglect the non-negativity constraints on c and ℓ for expositional purpose. The optimal consumption equals the inverse marginal utility of consumption of the shadow price of consumption,

$$c_t = u_c'^{(-1)}(\lambda), \tag{24}$$

and labor supply equals the inverse marginal utility of consumption of the negative after-tax shadow price of consumption,

$$\ell_t = u_l'^{(-1)}(-(1-\tau)\lambda). \tag{25}$$

The price of tomorrow's bond the household is willing to pay depends on the expected discounted shadow price of consumption tomorrow normalized by the shadow price of consumption today,

$$p_b = \beta \frac{\sum_{z^+} \lambda^+(z^+|z)\pi(z^+|z)}{\lambda}. \tag{26}$$

We set $\eta = 1$, $\sigma_1 = 1$ (log utility), $\sigma_2 = 1$ (log utility), and the discount factor $\beta = 0.96$ to facilitate comparisons with AMSS. We consider $\underline{c} \in \{0, 0.1\}$ where $\underline{c} = 0$ corresponds to AMSS. We show that $\underline{c} = 0$ has extreme and unnatural implications for tax revenues, whereas $\underline{c} = 0.1$ has more realistic implications. We set $\theta = 100$ which represents a situation where the penalty for missing the government spending target is high. The appendix discusses the implications of other utility function specifications.

We roughly follow Buera and Nicolini (2004) in their calibration of the transition probabilities to the US experiences in the last century: approximately two major wars per century with an average duration of 3 years. This implies the following transition matrix for the spending state:

$$\Pi = \begin{bmatrix} 0.9787 & 0.0213 \\ 0.3333 & 0.6667 \end{bmatrix}. \tag{27}$$

5.2 Feasible Region

The endogenously determined feasible regions of government debt are central to the debate of public debt levels. Thus, we present these feasible regions for the AMSS fixed-g and our flexible-g specification. They turn out to be fundamentally different: the fixed-g

parameterization implies almost no feasible government debt levels, whereas the flexible-g parameterization allows for significant levels of government debt.

For fixed-g, the feasible region depends on the level of government spending as well as the lower bar of consumption, $\underline{c}$. Figure 3 displays the feasible region in peace for three exogenous government spending processes ($g_{\text{peace}}, g_{\text{war}}$) being equal to (0.09, 0.27), (0.045, 0.135), and (0.0225, 0.0675). As one might expect, the feasible region expands for decreasing levels of exogenous government spending. In the high spending case—where spending equals 9% in peace and 27% in war of maximum GDP—there almost no feasible government spending. That is, the government can almost not be in debt to the households during peace. As we show later, there is no feasible positive debt level or war states in the high spending regime. For the two other cases with lower government spending, the feasible region extends to positive levels of debt such that the government can be in debt to the households.

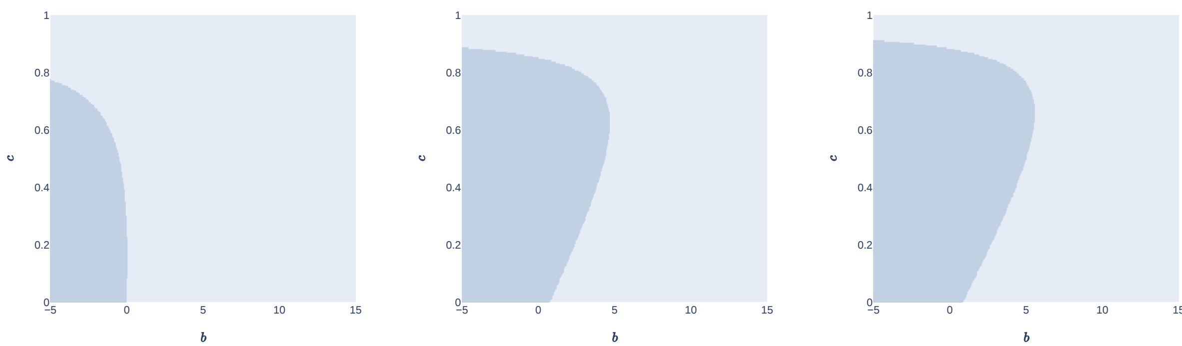

Figure 3: Numerically feasible region $\tilde{\Omega}(z = \text{peace})$ for fixed-g. The levels of government spending decrease from left to right: (0.09, 0.27), (0.045, 0.135), and (0.0225, 0.0675).

Figure 4 compares the feasible regions for $\underline{c} = 0.1$ (left) and $\underline{c} = 0$ (right), $\sigma_1 = \sigma_2 = 1$, and fixed-g spending of (0.09, 0.27). Specifying $\underline{c} = 0$ for the household's utility allows for government debt that is highly positive. For further results on feasible regions see Appendix A.1.

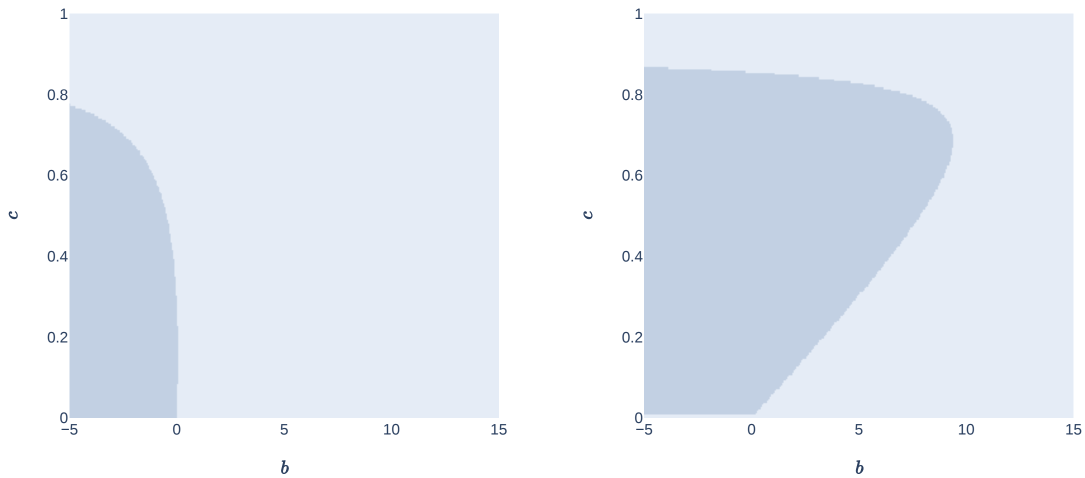

Figure 4: Numerically feasible region $\tilde{\Omega}(z = \text{peace})$ for fixed-g with $\underline{c} = 0.1$ (left) and $\underline{c} = 0$ (right).

We turn to the feasible region of the flex-g case illustrated in Figure 5. In contrast to the fixed-g case, the feasible region does not depend on the level of government spending and extends well into the positive b quadrant. Hence, it is feasible for the government to be in high debts to the households. The reason becomes clear from Figure 6: the government can always choose to spend less even though this is heavily penalized. Therefore, the feasible space includes all states such that it is possible to finance debt obligations while zeroing out any spending.

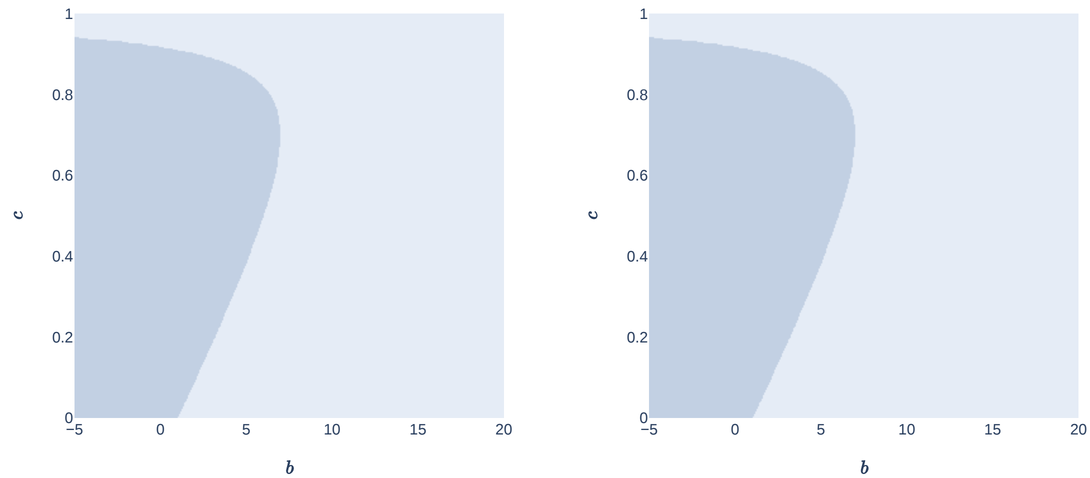

Figure 5: Numerically feasible region $\tilde{\Omega}(z = \text{peace})$ for flex-g with $\underline{c} = 0.1$ for high government spending $(0.09, 0.27)$ on the left and $(0.045, 0.135)$ on the right.

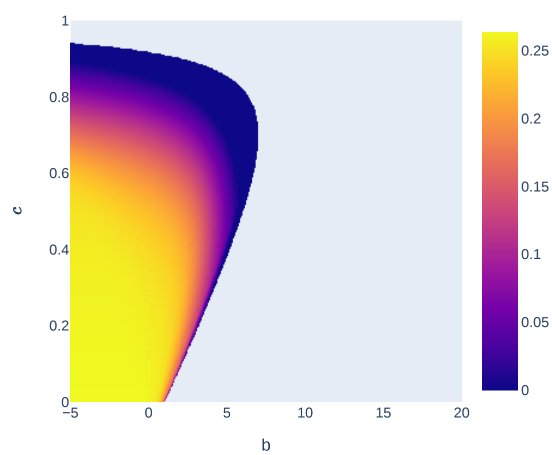

Figure 6: Government spending for flex-g in war.

Figure 7 shows the feasible region for changing lower bar of consumption $\underline{c} \in \{0, 0.1\}$. The lower bar of consumption $\underline{c}$ impacts the consumer similarly as in the fixed-g case. The feasible reason expands significantly for $\underline{c} = 0$.

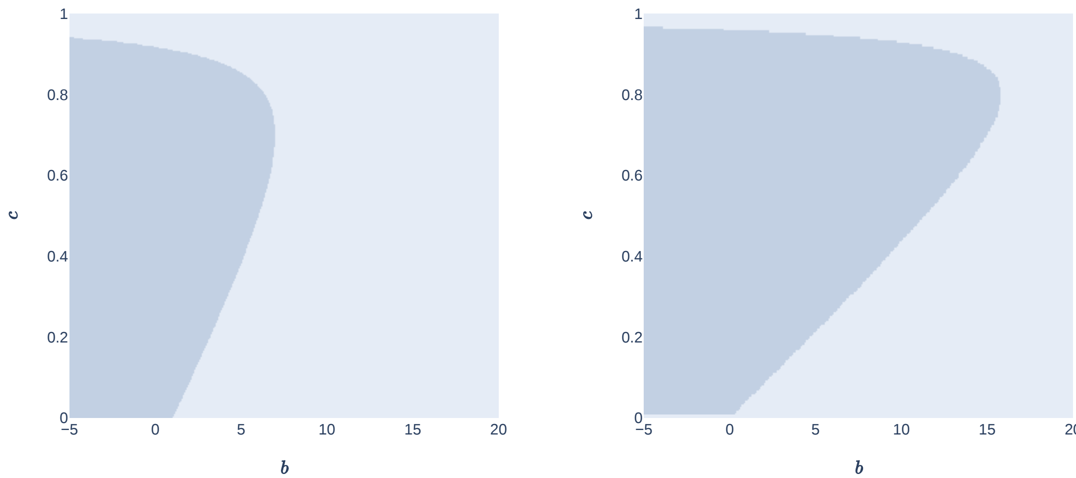

Figure 7: Numerically feasible region $\tilde{\Omega}(z = \text{peace})$ for flex-g with $\underline{c} = 0.1$ (left) and $\underline{c} = 0$ (right).

5.3 Endogenous Government Spending: Global Dynamics

Figure 8 displays how the (b, λ) state changes in all four contingencies: in peace and transition to peace, etc. The key result here is that most of the feasible region is comprised of transient states. If $\lambda > 5$ then λ in the next period is below 5. Similarly, we see that the lowest λ states are also transient. Similarly, the next period's debt level is less than 4 no matter what the current debt level is. We have imposed the constraint that $b > -3$, but that is without loss of generality. If b=-3, we are in a war-chest state. Any further reduction in debt (that is, an increase in government assets) would only produce more interest income that is then paid back to agents through a larger transfer. In this case, we can conclude that the only states that are visited after the first period are in the box defined by $-3 < b < 3.12826$ and $1.93486 < \lambda < 4.38104$. We call that the "ergodic box" because it contains all non-transient states. Our simulations below will be confined to that box.

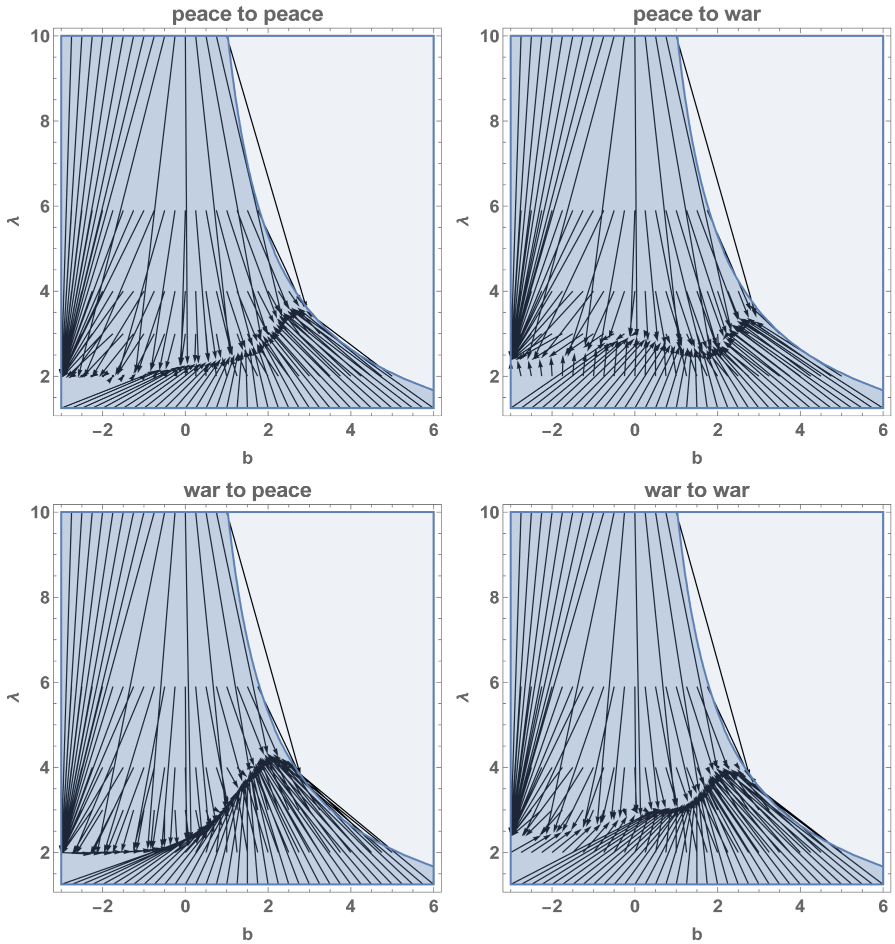

Figure 8: One-step transitions for (b, λ).

5.3.1 Permanent Peace versus Permanent War

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 display the paths followed if the economy begins with $(b, \lambda) = (b_{\text{init}}, \lambda_{\text{init}})$ and proceeds in perpetual peace or war for 400 periods.

In perpetual peace, both the b and λ paths move to their long-run values as depicted in Figure 9. For almost all fixed initial λ and initial b, the government pays off its debt and converges to some level of positive government assets which depends on the initial λ. Even though the government builds up assets, it only builds up a "war chest" if it starts with a high initial endowment of assets. There is no drift to high debt.

The perpetual war case in Figure 10 is different. For most initial levels of debt, there is a gradual climb in debt. The intuition is clear: When in war, there is a 1/3 chance of peace in

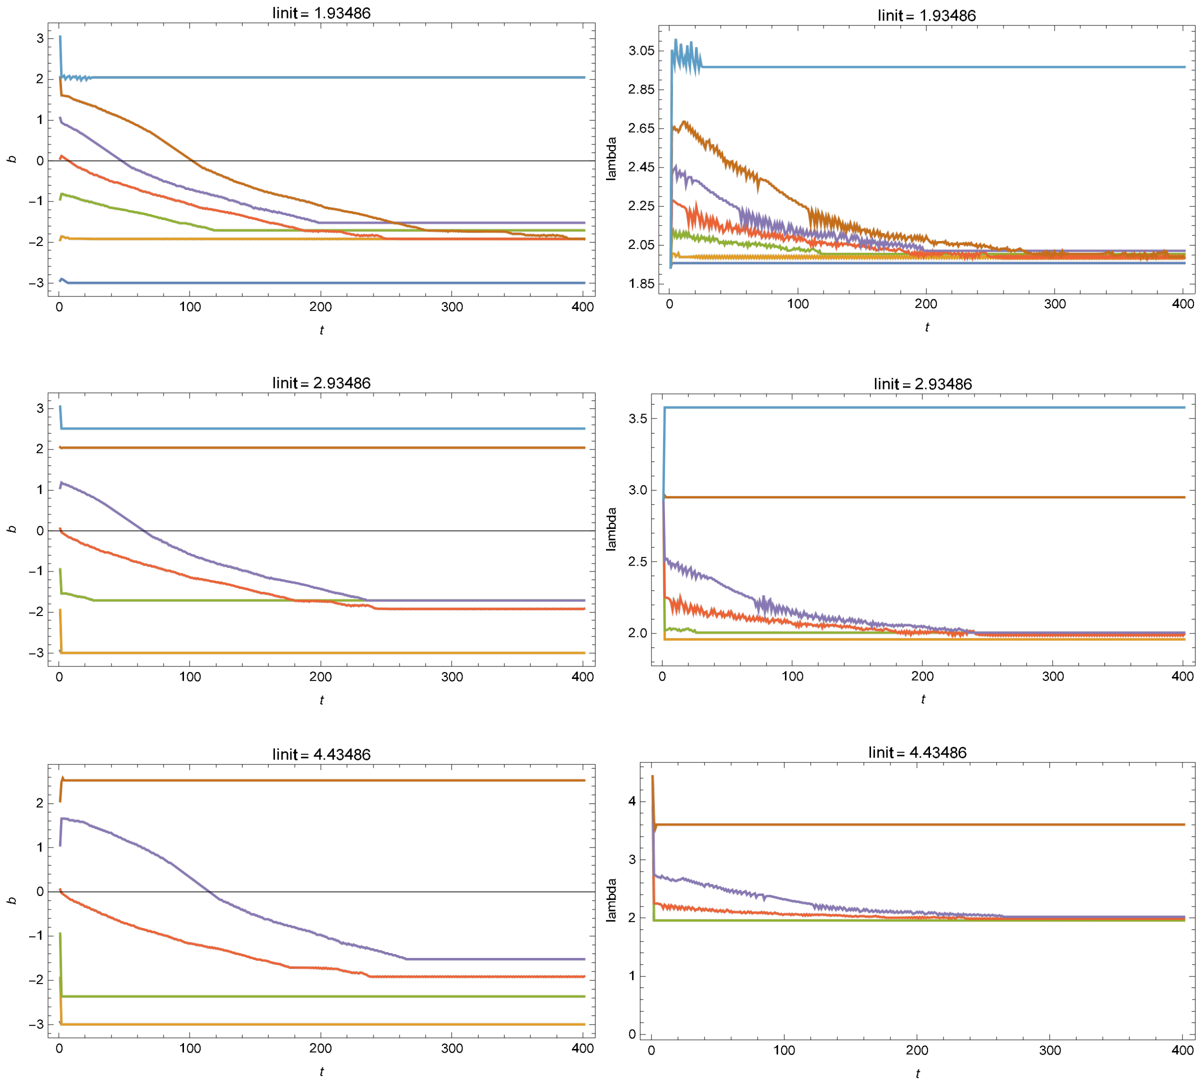

Figure 9: Perpetual peace time-series evolution for b (left) and λ (right). From top to bottom, λ_{init} equals 1.93486, 2.93486, and 4.43486.

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. In this situation, the government cuts back on its expenditures. There exist some starting configurations which end up in a war chest.

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, but they may also seem counter-intuitive at first. Perpetual wars do not seem to wipe out war chests

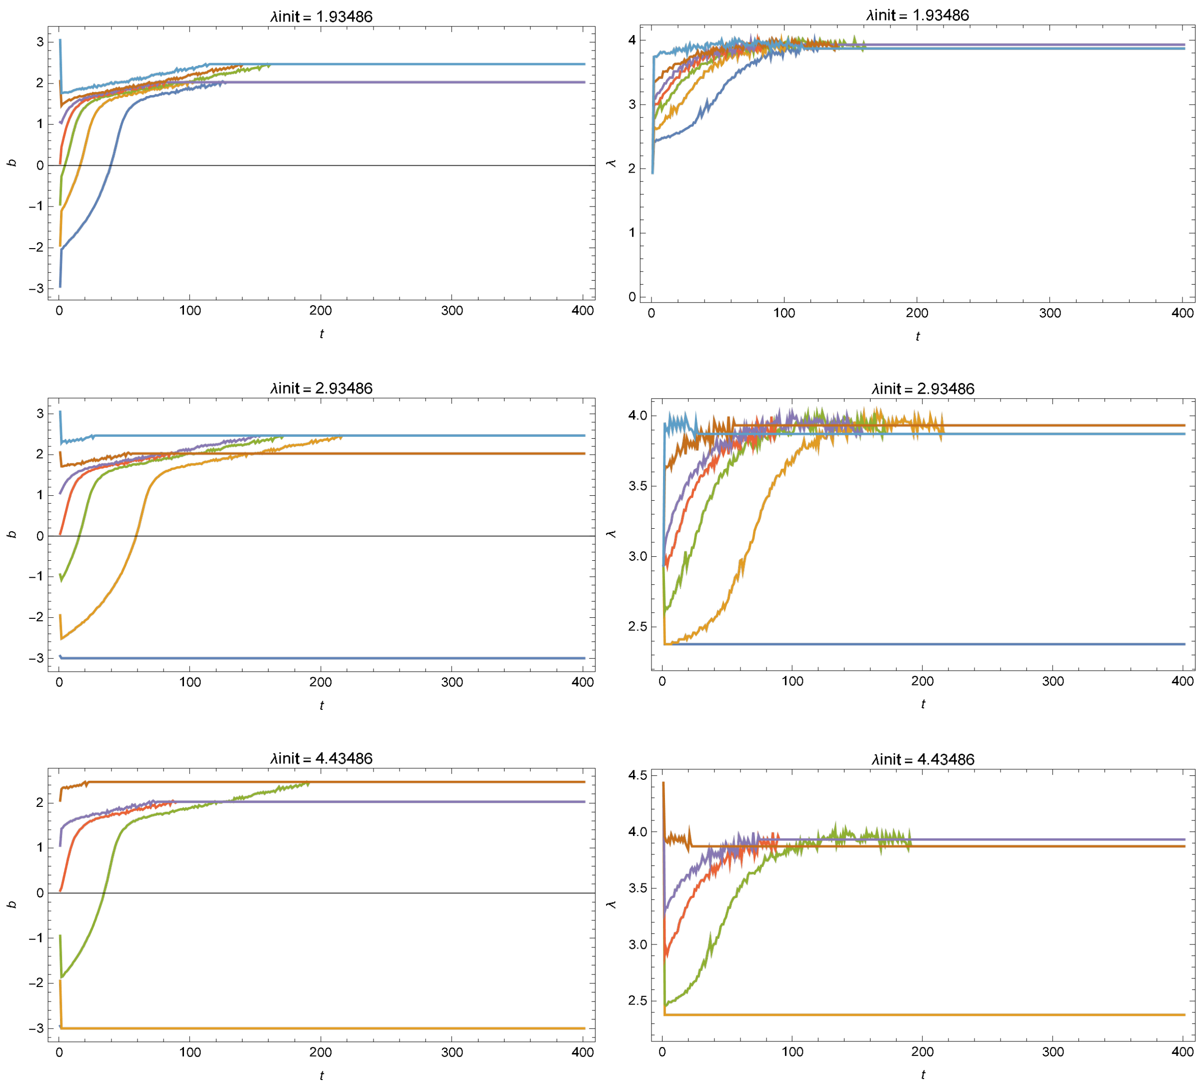

Figure 10: Perpetual war time-series evolution for b (left) and λ (right). From top to bottom, λ_{init} equals 1.93486, 2.93486, and 4.43486.

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.

5.3.2 Random Paths

We next analyze the movement of debt, b, the shadow price of consumption, λ, the government spending, $g(z)$, and the tax-rate, τ, for a stochastic path of peace and war, generated according to the Markov transition matrix Π.

In Figure 11, the government starts with ($b_{init} = 1.0, \lambda_{init} = 2.4$), i.e., the government starts in debt. After a peak in debt, the government assets increase to $b = -1.5$, which the government then mostly sustains for the entire horizon. The marginal utility of consumption, λ, jumps between 2 and 2.6 for peace and war, respectively. The government approximately meets target spending $\bar{g}(z)$ for peace and war. In peace, the tax rate of about 0.05 is relatively low and is raised up to 0.35 to finance war spending. This appears to violate the Barro argument that tax rates should be smoother than spending. This is where the linear-quadratic approach misses some important microeconomic considerations. Optimal tax rates depend on elasticities, and the fluctuations in consumption necessitated by the fluctuations in government spending will affect the elasticities of labor supply. This shows that even when the only asset is safe debt, elasticities will play an important role in labor tax rates. We expect that these features will be sensitive to assumptions about labor supply and that there may not be any robust results when we examine a range of elasticities implied by empirical work.

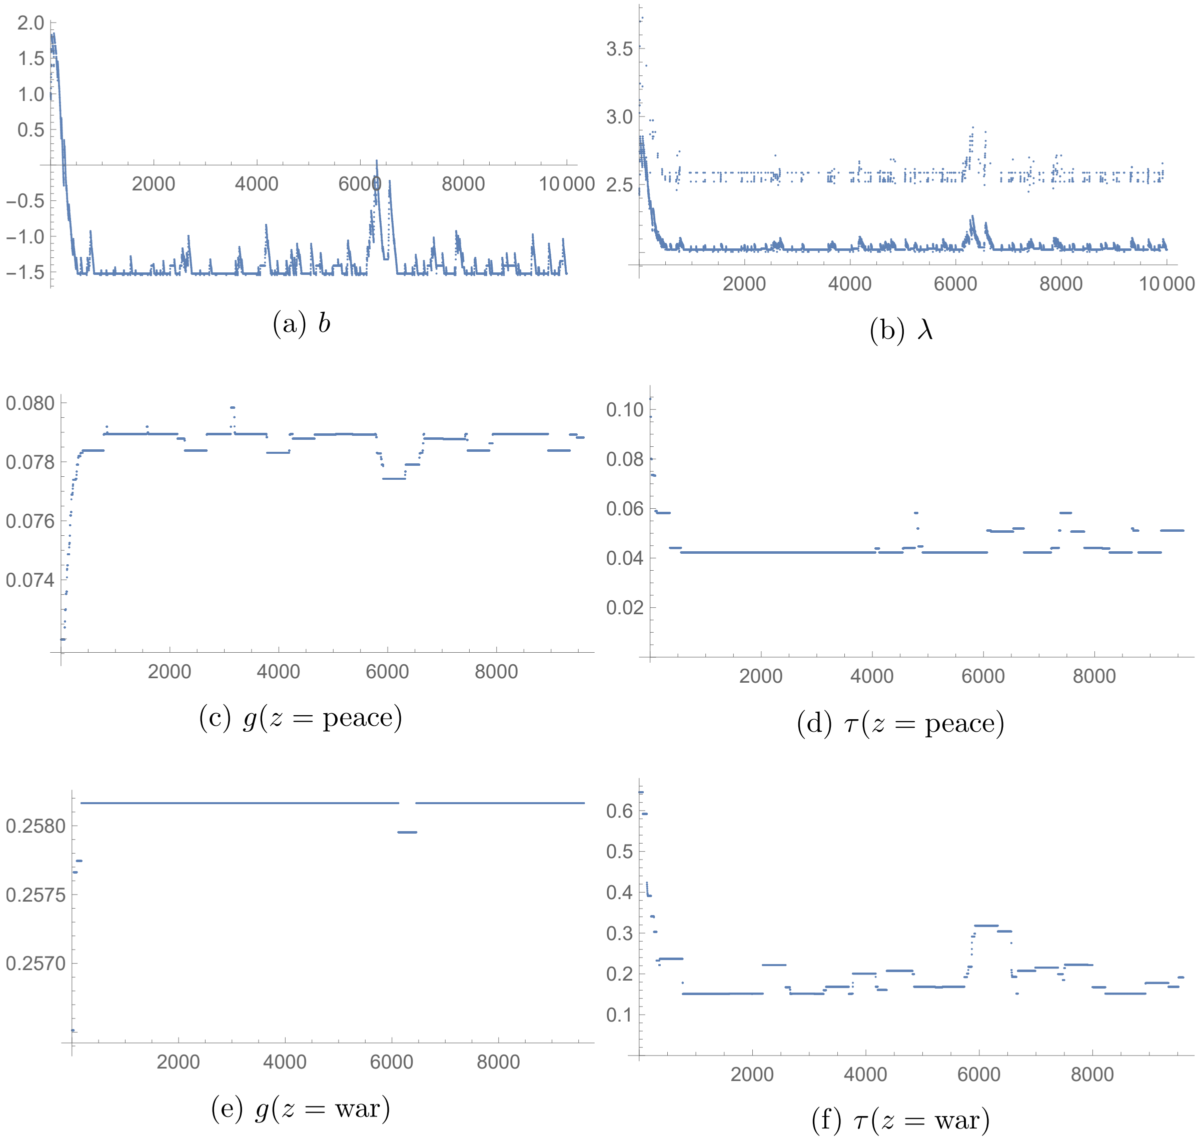

(a) b (b) λ

(c) $g(z = \text{peace})$ (d) $\tau(z = \text{peace})$

(e) $g(z = \text{war})$ (f) $\tau(z = \text{war})$

Figure 11: Evolution of b, λ, $g(z)$ and $\tau(z)$ for a stochastic peace/war process. We calculate the peace (war) values as the minimum (maximum) over a 200 year window.

In contrast, the government follows a different policy path when starting at a higher debt level as depicted in Figure 12. The government starts with ($b_{init} = 2.0, \lambda_{init} = 2.4$). Instead of paying off the debt, the government moves towards a higher debt level of $b = 2.85$. To sustain this government debt level, the government cuts down on government spending in peace and in war down to almost 0 and 0.14, respectively. To finance the debt, the tax rates stay at a high level of about 0.7 in war and 0.5 in peace.

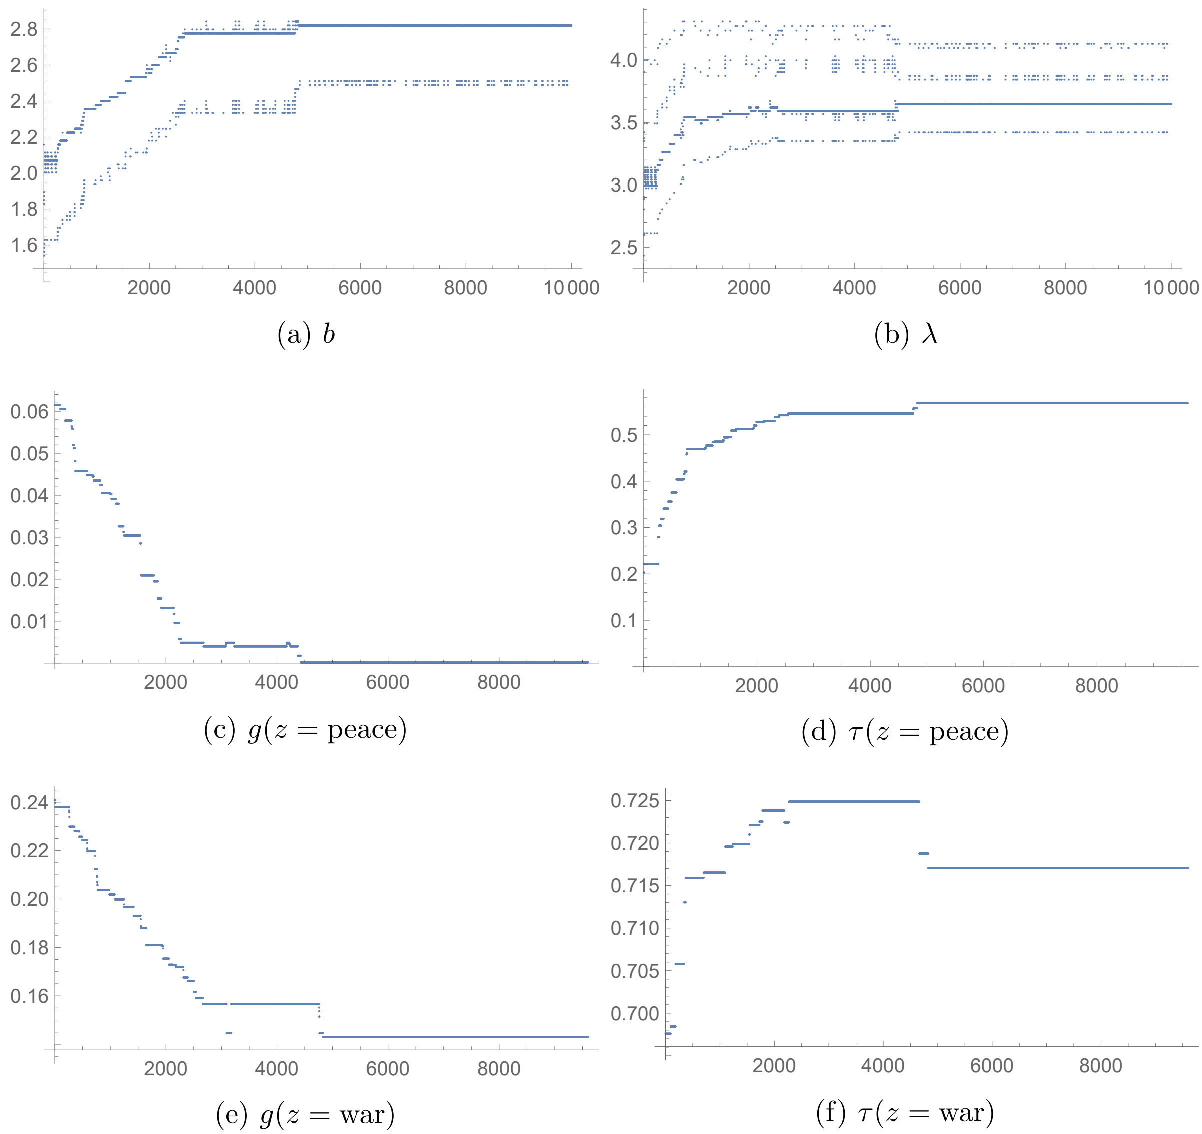

Figure 12: Evolution of b, λ, $g(z)$ and $\tau(z)$ for a stochastic peace/war process. We calculate the peace (war) values as the minimum (maximum) over a 200 year window.

5.3.3 Long-run States

Next, we analyze the co-movement of (b, λ). Figures 13 and 14 visualize all states that are visited. The first plot shows the (b, λ) states visited for the first two periods, as indicated by the label $1 \leq t \leq 2$. Analogously, the three other plots display the points visited over some later period of time. These resemble the same long-run dynamics visualized in the previous section; we intentionally omit cases that start from points with a very low b, as they quickly converge to a war chest and stay there.

The policy shown in Figure 13 corresponds to the long-run time-series depicted in Figure 11. The government starts in a low debt region, and initially, moves towards a higher debt level. After $t = 100$, the government builds up assets and moves to $b < 0$. After $t = 1000$, the policy moves between two distinct sets where the government has positive assets.

Figure 14 is typical if the initial debt higher. The dynamics correspond to the time-series plots in 12. There is a slow drift towards high debt. In fact, the last 4000 periods are spent close to the boundary of the feasible region. 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.

5.4 Lessons from our analysis

While we have discussed the results from just one example, they do point to several basic insights that are unlikely to be reversed by alternative parameter choices. First, the feasible level of debt depends greatly on the flexibility of government spending levels. The assumption of fixed spending levels imposed by exogenous factors makes it very difficult to support debt and, in cases similar to US spending levels, even impossible to have any positive level of debt.

Over the long-run, the economy could get stuck with a high level of debt and taxes but low levels of government spending, consumption, and labor supply. While these are all feasible long-run conditions with respect to the economic conditions we model, there may be a question of whether such a condition is politically feasible. In our model, utility over government spending is separable. Our model has the alternative interpretation is that only the decision-maker cares about government spending. In that case, one could alter the utility over g to represent fear on the part of government officials of losing power, being tarred and feathered, or other undesirable reactions from the public. We leave further consideration of this for future work.

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 levels of government spending. In fact, it is easy to find cases where no positive initial level of debt is feasible.

Our analysis assumes that government expenditure is endogenous and chosen according to a utility function that is subject to taste shocks. 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

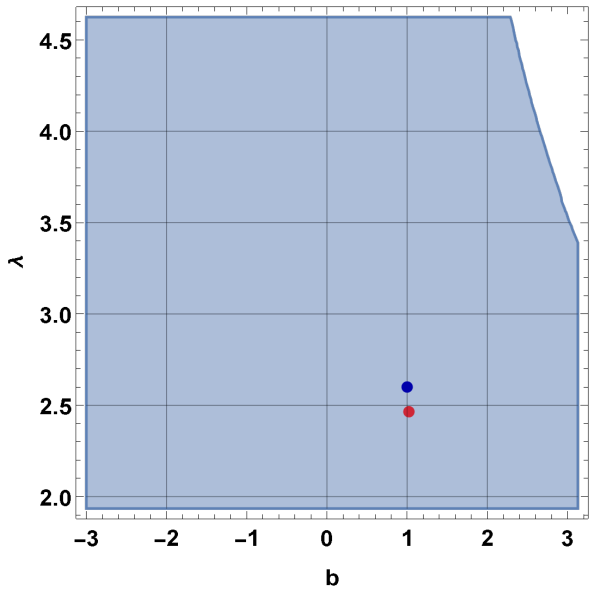

(a) $1 \leq t \leq 2$

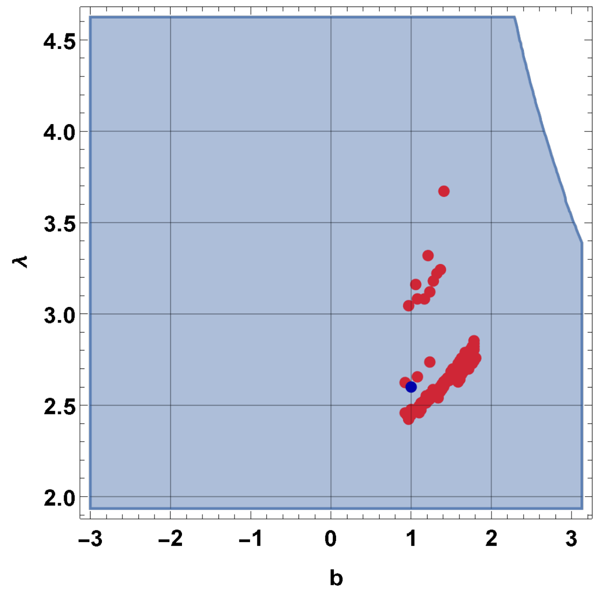

(b) $2 < t \leq 100$

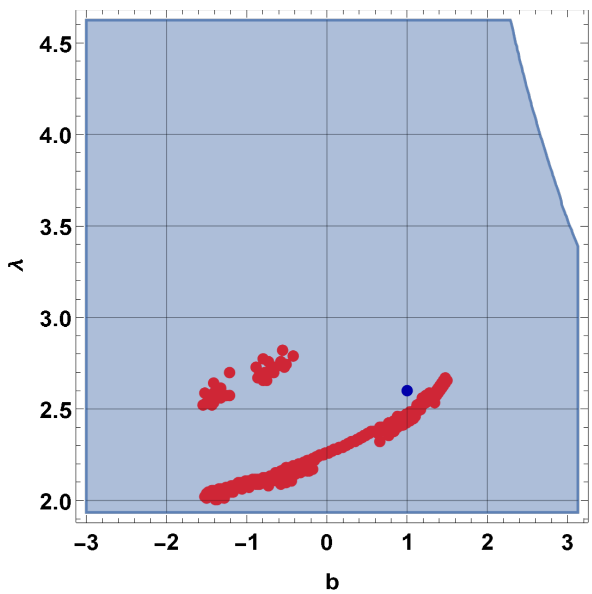

(c) $100 < t \leq 1000$

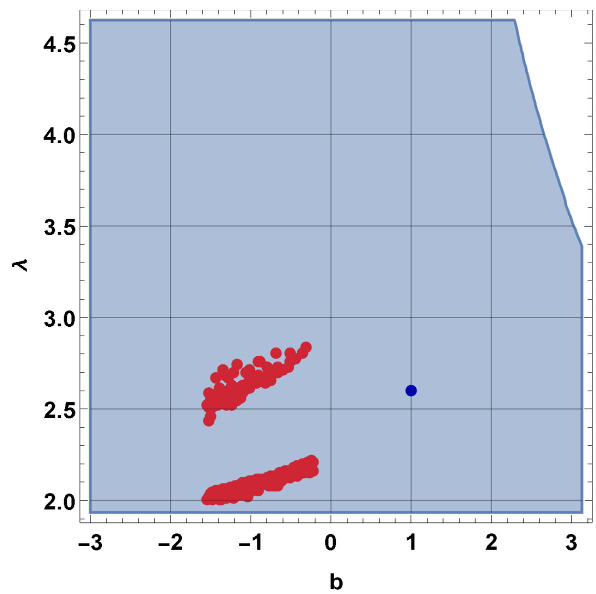

(d) $1000 < t \leq 5000$

Figure 13: All states that are visited (red dot) starting from an initial state (blue dot). From left to right, top to bottom, states visited between $1 \leq t \leq 2$, $2 < t \leq 100$, $100 < t \leq 1000$, and $1000 < t \leq 5000$.

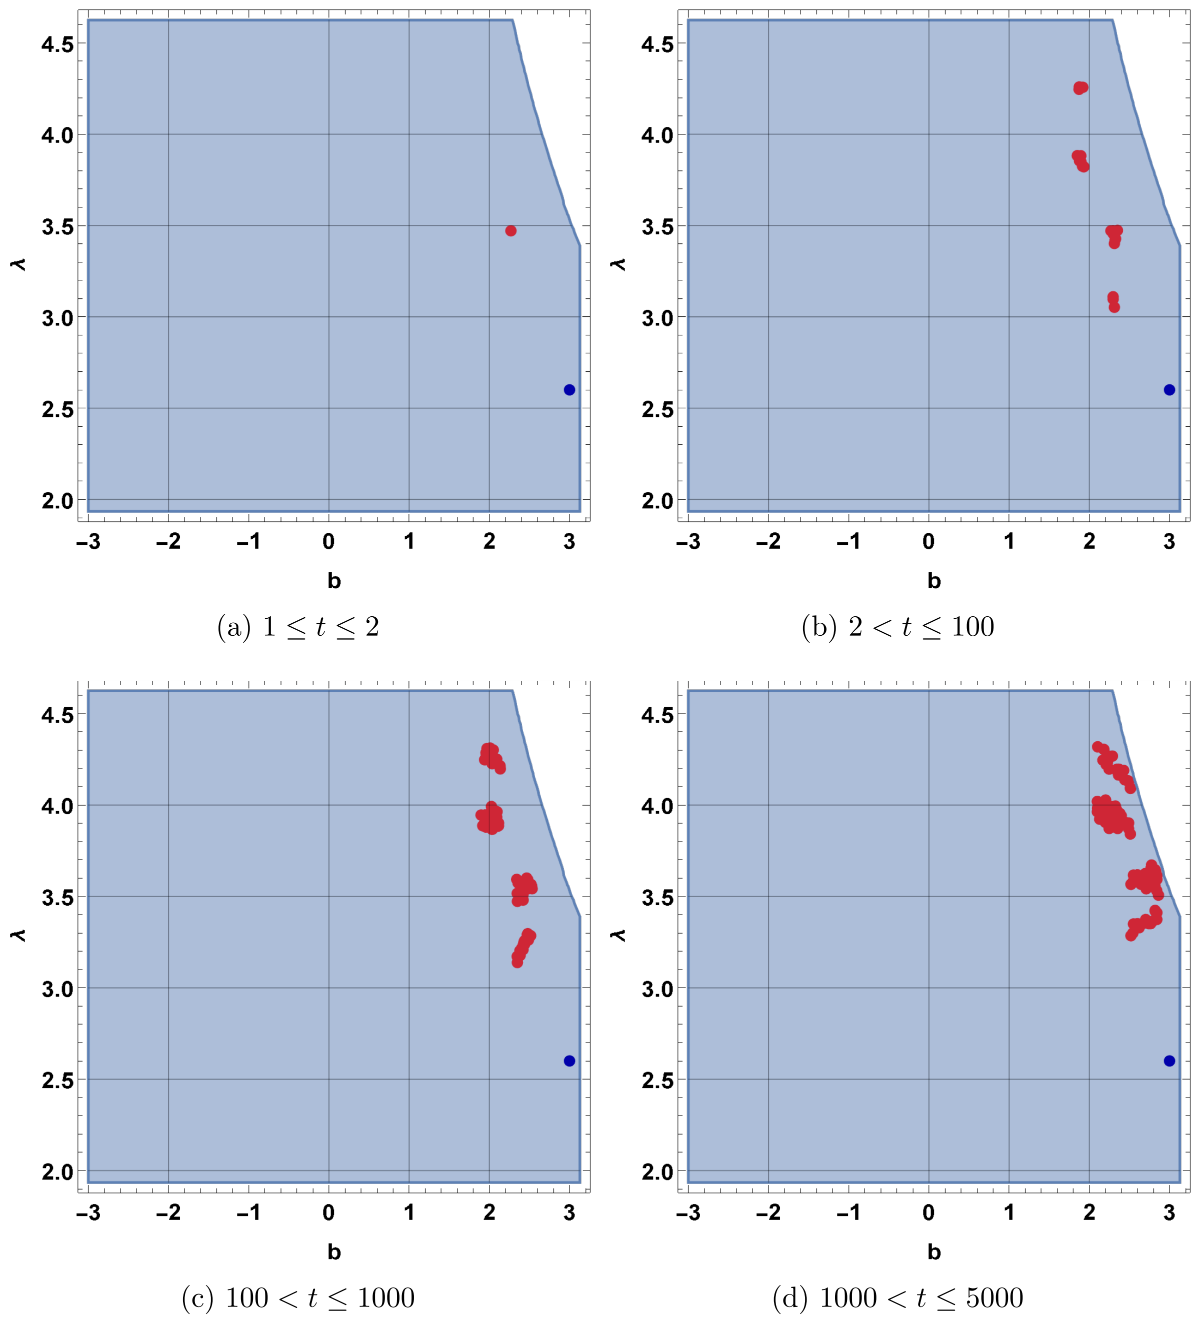

Figure 14: All states that are visited (red dot) starting from an initial state (blue dot). From left to right, top to bottom, states visited between $1 \leq t \leq 2$, $2 < t \leq 100$, $100 < t \leq 1000$, and $1000 < t \leq 5000$.

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.

Models, like AMSS and ours, that assume one only factor (labor), one consumption good, and identical agents are highly stylized. While we argue that some issues, like the importance of endogenous spending, can be discussed in such a model, many other issues cannot. For example, the absence of physical capital allows interest rates to fluctuate far more than is empirically reasonable. There are some aspects of our results, such as the initial response to a war shock, which we did not discuss because they depend on this excessive ability of the interest to change. Some aspects of our computational analysis are fancier than necessary for this problem but we are looking forward to solving models with capital where the curse of dimensionality will start to have a bite if we have to rely on three-dimensional discretizations of the three-dimensional state space.

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 on the web), we will be able to explore alternative hypotheses based on the data provided by the library of simulations.

References

Aiyagari, S. R., 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

Bizer, D. S., & Durlauf, S. N. (1990). Testing the positive theory of government finance. *Journal of Monetary Economics*, *26*(1), 123–141. doi: 10.1016/0304-3932(90)90034-2

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

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

Marcet, A. (1988). Solving Non-linear Stochastic Models by Parameterizing Expectations. *Manuscript. Pittsburgh: Carnegie Mellon Univ.*.

A Appendix

A.1 Exogenous Government Spending

This section presents additional results for the exogenous government spending case—the standard AMSS economy. Specifically, we look at the cases presented in Table 1, where we reduce the government spending successively to study the resulting feasible regions.

Case #	σ_1	σ_2	$\underline{c}$	exogenous g	$\bar{g}$	c bounds
1	1.0	1.0	0	Yes	$\{0.09, 0.27\}$	[0.01, 1]
2	1.0	1.0	0	Yes	$\{0.045, 0.135\}$	[0.01, 1]
3	1.0	1.0	0	Yes	$\{0.0225, 0.0675\}$	[0.01, 1]

Table 1: Experiment setup

We study the feasible region for cases 1-3. Compared to the setting before, we additionally change Π to 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 1b denotes case 1 with the symmetric matrix with 66.67% "staying-where-you-are" probability. Figure 15 depicts the feasible regions.

Cases 1(a-c) and 2(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 between a, b, and c do not qualitatively change this result. For very low government spending—in peace as well as in war—as in case 3(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

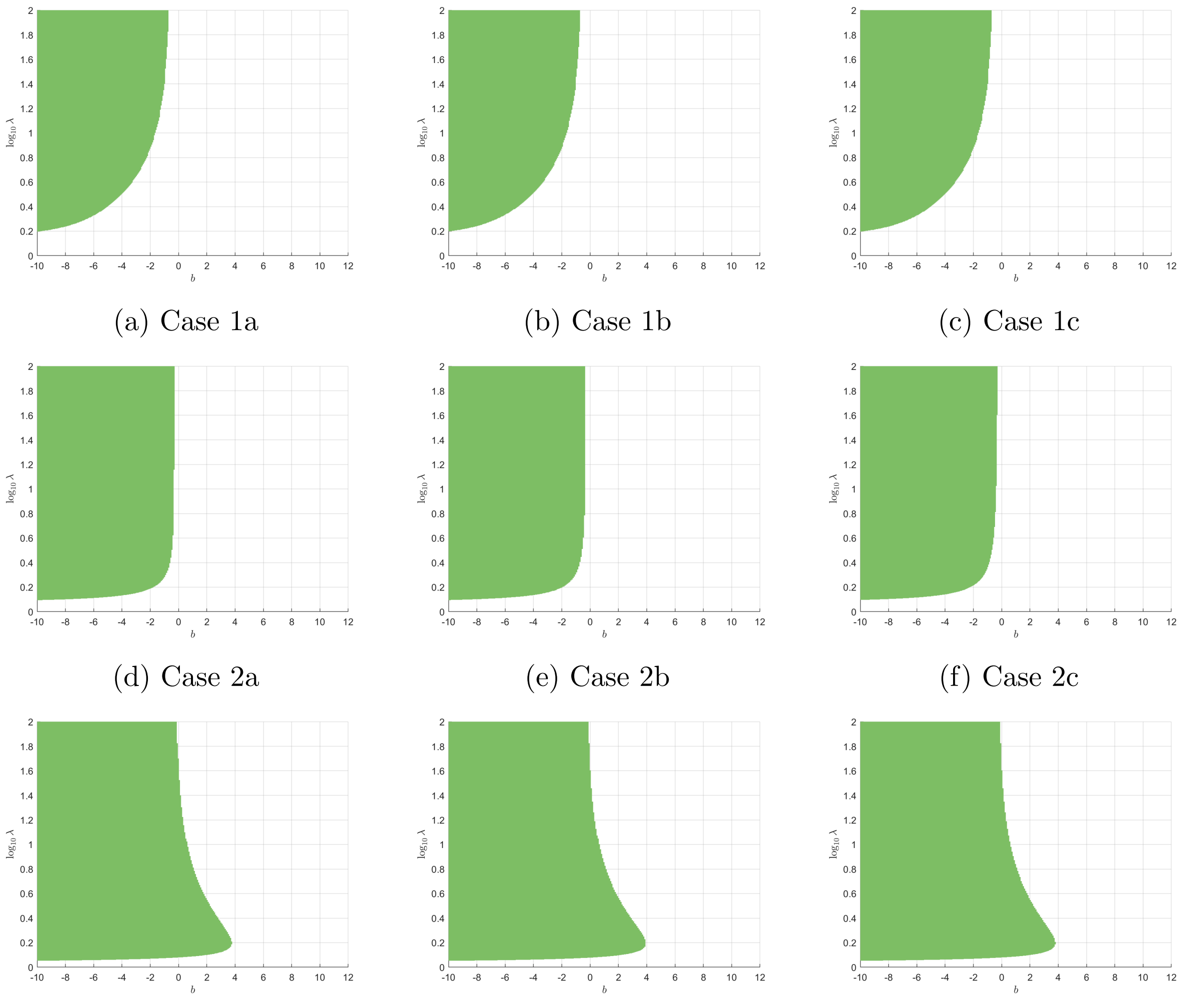

(a) Case 1a (b) Case 1b (c) Case 1c

(d) Case 2a (e) Case 2b (f) Case 2c

(g) Case 3a (h) Case 3b (i) Case 3c

Figure 15: Numerically feasible regions $\tilde{\Omega}(z)$. Top to bottom, cases in Table 1. Left to right, symmetric transition probability matrices Π with a 50% (a), 66.67% (b) and 80% (c) probability of "staying-where-you-are".

computation, the cases, and their implied feasible regions. In later versions of this paper, we analyze the time-series simulations for these cases as well.

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 individual optimization problems. As optimizer, we use IPOPT by **?** in combination with the automatic differentiation tool CasADi by Andersson, Gillis, Horn, Rawlings, and Diehl (2018). The code is 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 2. Recall the utility function from Equation **??**

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

Case #	σ_1	σ_2	$\underline{c}$	exogenous g	$\bar{g}$	b bounds	c bounds
4	0.5	0.5	0.1	No	$\{0.09, 0.27\}$	$[-10, 12]$	[0.01, 1]
5	0.5	1.0	0.1	No	$\{0.09, 0.27\}$	$[-10, 12]$	[0.01, 1]
6	1.0	0.5	0.1	No	$\{0.09, 0.27\}$	$[-10, 12]$	[0.01, 1]

Table 2: Experiment setup

The utility functions are now combinations of sqrt utility functions and log utility functions. Figure 16 depicts the feasible regions in this case. The general feasible debt level is in cases 4 and 5 lower than before, but qualitatively, they show the same pattern as before: The level of feasible debt shrinks with increasing λ.

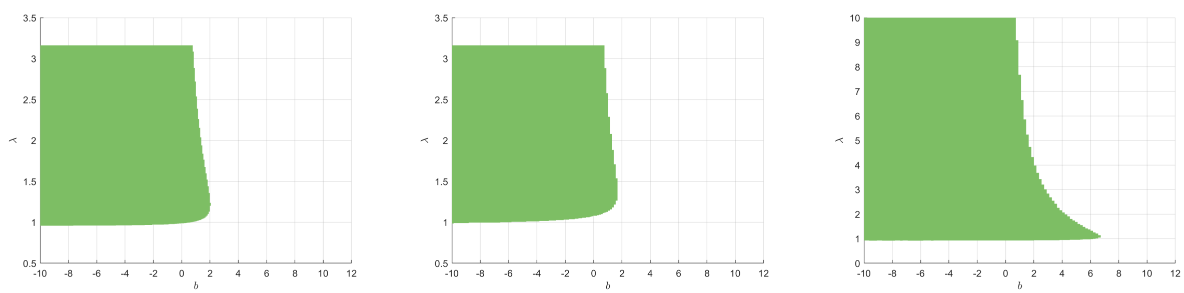

(a) Case 4 (b) Case 5 (c) Case 6

Figure 16: Numerically feasible regions $\tilde{\Omega}(z)$ for cases in Table 2

A.3 Comments on the AMSS Algorithm

We examined the sensitivity of the AMSS code to changes in the initial debt level and the seed in the random number generator.

Table 3 displays how the initial debt level used in the simulations affected the coefficients in the solution. We see that different initial debt levels affect the coefficients in the solution.

Table 4 displays how the seed used in the simulations affected the coefficients in the solution. The AMSS code set the seed to equal 1. We found that changing the seed to 2 affected the coefficients in the solution, and that the algorithm did not converge for seeds 3, 4 and 5.

The sensitivity of results to the seed is common in methods that rely on Monte Carlo simulations. Any use of Monte Carlo should be repeated enough times with different seeds

Initial debt	0	50	100	200	400
β_1^1	-6.870E+00	-6.873E+00	-6.875E+00	-6.877E+00	-6.877E+00
β_2^1	2.992E-01	2.952E-01	2.915E-01	2.861E-01	2.767E-01
β_3^1	8.993E-02	9.498E-02	9.990E-02	1.086E-01	1.380E-01
β_1^2	-2.099E+00	-2.099E+00	-2.099E+00	-2.099E+00	-2.098E+00
β_2^2	6.704E-04	7.106E-04	7.218E-04	6.817E-04	8.697E-04
β_3^2	-3.736E-04	-3.580E-04	-3.420E-04	-3.127E-04	-2.288E-04

Table 3: The parameters of the first parameterized expectation, $\{\beta_i^1\}_{i=1}^3$, and of the second parameterized expectation, $\{\beta_i^2\}_{i=1}^3$, for changing levels of initial debt.

Seeds	1	2	3	4	5
β_1^1	-6.870E+00	-6.867E+00	-	-	-
β_2^1	2.992E-01	3.018E-01	-	-	-
β_3^1	8.993E-02	8.981E-02	-	-	-
β_1^2	-2.099E+00	-2.098E+00	-	-	-
β_2^2	6.704E-04	7.863E-04	-	-	-
β_3^2	-3.736E-04	1.995E-04	-	-	-

Table 4: The parameters of the first parameterized expectation, $\{\beta_i^1\}_{i=1}^3$, and of the second parameterized expectation, $\{\beta_i^2\}_{i=1}^3$, for changing seeds.

to compute the variance of the solution. If that variance is ”small” then on could accept any one of the solutions. For any particular application, one should define ”small” in some economically relevant fashion as well as demonstrate that the variance is ”small”. In the case of AMSS, we will (in another paper) examine how different solutions affect time series paths and other economically important implications of the various solutions.