Optimal Dynamic Stochastic Fiscal Policy with Endogenous Debt Limits

Kenneth L. Judd* Alice Feng† Şevin Yeltekin‡

Philipp Müller §

October 16, 2024

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 Barro (1979) and Aiyagari et al. (2002) dynamic fiscal policy models with incomplete risk markets but introduce essential changes. First, we assume that government spending is endogenous, chosen according to a utility function. Second, we formulate the government's dynamic incentive problem as dynamic programming. 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 general tendency to accumulate a war chest large enough to allow taxation to disappear, even during long periods of peace. Second, flexibility in government spending substantially increases the capacity to issue debt. If a country can adjust

*Hoover Institution & NBER, kennethjudd@mac.com

†Hoover Institution

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

§Accenture AG, Switzerland, philipp.mueller@business.uzh.ch. Philipp Müller gratefully acknowledges the support of Kenneth Judd, Senior Fellow at the Hoover Institution.

its spending when faced with financial distress, it can sustain a credible reputation for honoring high debt levels. Third, the debt and tax processes are not stationary in any useful sense.

1 Introduction

US public debt levels and deficits have reached record levels in the past 20 years due to a combination of war, tax cuts, spending increases, financial crises, and a pandemic. This alarming development plus projections of future spending due to the aging of the US creates major concerns about the fiscal system's ability to finance government expenditures. Central to these concerns are the optimal tax policy and the maximum level of public debt that fiscal policy can sustain in the long run—the endogenous debt limit.

Barro (1979) develops a positive theory of government finance. In his model, the government faces an exogenous government expenditure process financed by a combination of taxation and the sale of one-period, risk-free bonds. The economic costs of taxation (inefficiencies, enforcement costs, etc.) are convex and increasing with the level of taxes. Barro (1979) argues that debt allows the government to smooth taxation to reduce the costs of taxation. He then argues that this implies random walk behavior for taxes. While this is true if the costs of taxation are quadratic (an assumption that violates any model of taxation consistent with a finite budget constraint for agents), we will show below that this is not true in general.

In contrast, Lucas and Stokey (1983) examine similar issues in a model with complete markets. Complete markets imply that the value of government debt depends on the state. Therefore, a bond purchased today will have value tomorrow that depends on tomorrow's contingencies, such as productivity and the taste for government spending. The tax rates depend on the movement of elasticities, and state-contingent debt will absorb shocks to spending and revenues. The shadow price of government revenue will be constant across states of nature, and tax rates in any period are governed by conventional optimal-tax rules that focus on elasticities. Lucas and Stokey (1983) is essentially a dynamic example of the Diamond-Mirrlees theory of optimal taxation. In Barro (1979), the tax policy inherits the serial correlation patterns of the government spending process and the evolution of debt. This happens in the Lucas-Stokey model only to the extent that elasticities that determine tax rates are correlated with spending and debt, something which may be true in special cases but not in general.

Aiyagari, Marcet, Sargent, and Seppala (2002) (hereafter referred to as MSS [1]) attempts to examine a full rational expectations, general equilibrium model, as was Lucas-Stokey, but restrict government debt to be safe as in Barro (1979). While real debt is not perfectly safe, the debt for the US and the UK are relatively safe. The absence of state-contingent debt constrains the government's ability to smooth taxes over states. MSS's suspension of complete markets aims to rationalize the outcomes Barro had predicted.

In MSS, rational economic agents make consumption, labor supply, and asset accumulation choices. They show that if agents are risk neutral (the Aiyagari suggestion) the government builds up a war chest, that is, the government accumulates enough assets to set the tax rate to zero and finance all spending out of its asset income. Their analysis with risk averse agents argues that taxes will follow martingale behavior similar to claims in Barro (1979). However, they use a parameterized expectations algorithm (PEA) by Marcet (1988) to compute solutions to those cases. We have examined their code and find it to be unreliable in solving the MSS model in many ways: they assume a special functional form for policy functions, results depend on the seed of a random number generator, they make no effort to verify the accuracy of their solutions, and they produce results very different from our – verified – results based on standard numerical techniques. We believe that the main problem is that PEA only aims to make

[1] We find many problems with this paper. Earlier versions did not include Aiyagari as an author. Aiyagari was a referee who suggested that they examine the case of risk neutral agents. This case appeared in the published paper. After Aiyagari's death, the authors added his name to the paper in recognition of his contribution. Aiyagari's example is the only part of the paper that we do not criticize. We believe that it is inappropriate to imply that any errors were due to Aiyagari. The previous working paper versions were MSS papers and include useful details about the analysis in the published paper. Therefore, we will use the abbreviation MSS.

the L2 norm of Euler equation errors small in a simulated time series path. The L2 norm allows for large errors at many points in the state space, errors which may significantly affect outcomes. We will use dynamic programming methods and impose optimality conditions at every state.

MSS assumes exogenous government spending, the typical assumption in the optimal tax and debt policy literature. We assume that spending is determined by a social utility function and its shocks. We claim this is a more realistic treatment of history, particularly for the US and UK in the 20th century. We find that the fixed-g and flexible-g models produce dramatically different results, particularly when we calibrate government spending shocks to the US experience in the last century. Exogenous spending is often excused as a simplification, but we show that it is a critical feature of any model. We will use the same micro foundation model as MSS but we will allow flexible government spending. 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 MSS fixed-spending model and find positive persistent feasible levels of debt. We show that flexibility in spending allows persistent high levels of debt that look like historical experiences. The difference is that if hit with an unlikely persistent sequence of positive spending shocks, we assume that the government will reduce spending as it approaches infeasible levels of debt, a response not allowed in the MSS analysis. The conclusion is obvious: feasibility of debt depends critically on the ability to remain solvent even when faced with a low probability sequence of adverse events.

MSS uses the "first-order approach" as defined in Marcet-Marimon to formulate the government's optimization problem. Unfortunately, this approach is not reliable and the analysis in MSS fails to identify the set of feasible debt levels. We instead formulate the policy problem as an infinite horizon dynamic programming problem following Kydland and Prescott (1980).[2] They emphasize that a key challenge for such problems is the endogenous nature of the feasible set of states, defined as the set of states for which there exists a government policy sequence such

[2] **?** claimed dynamic optimization methods could not be applied to economic planning problems, an assertion which was contradicted by many papers published before 1977.

that no equilibrium condition or budget constraint is violated. Determining the feasible region introduces multiple challenges: (i) we need to determine it jointly with solving the dynamic programming problem, and (ii) the government's value function diverges to $-\infty$ close to its border, a fact making it difficult to approximate the value function.

In this paper, we use discrete state dynamic programming, which avoids the necessity of specifying an approximation method for the value function. Dynamic programming involves solving a Bellman equation at each state. Since we cannot prove that the value function is concave, we have to treat each Bellman equation as a global optimization problem. In fact, we compute all feasible choices. We use the computational power of graphical processing units (GPUs) to perform the heavy computational burden. While taxation and debt are the focus of this application, it is clear that our computational framework can solve a wide variety of recursive contract problems.

We compute the optimal tax and debt policy assuming endogenous government spending using our computational framework. Our analysis focuses on the distribution of debt over time as well as the endogenous feasible region implied by the assumptions on the primitives of the economic environment. Our results have substantially different implications with much more complex policies than earlier analyses implied. In particular, we find that the stochastic process of debt is far more complex than the simple random walk claim.

Further, the complexity of the tax and debt process suggests that it cannot be captured well by any simple auto-regressive process. War chest accumulation highly depends on the initial government assets: if the initial government assets are high enough, the government accumulates a war chest which allows long-run tax rates to be zero. If the government assets are less than a war chest, governments react to transitory tastes for high spending by issuing debt and raising taxes. Governments may reduce their debt during extended periods of low spending, but that is ultimately overcome by occasional periods of a taste for high spending.

Our model shares two key assumptions with Barro (1979) and MSS(2002).

First, 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. Both the US and the UK have used inflation to default on debt partially, 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. It may be true that there is significant default risk for the debt of most countries, but a large fraction of sovereign debt is safe because of the large market share of the UK and US debt plus that of other safe countries. The presence of default risk will likely reduce the feasible levels of debt, but the impact will be sensitive to the details of the default process. We stay with the assumption that government debt is risk-free because it allows us to compare our results to other work, and it appears to be a reasonable first-order assumption to make for the US, UK, and some other major economies. Second, we assume precommitment by the government to its tax policy. Therefore we do not discuss time-consistency issues of the optimal tax policy.

Optimal macroeconomic policy problems design typically involve solving high-dimensional dynamic models with parameter values advocated by macroeconomists. Examining the uncertainty due to parameter uncertainty (an activity which is called Uncertainty Quantification) showed that the current social cost of capital is uncertain but contained in the 40-100 dollar per ton range, substantially smaller than what many advocate. The analysis in Cai et al. (2017) and Cai and Lontzek (2019) required the use of a supercomputer (Blue Waters). It is quite reasonable to question the accuracy of any such complex and massive computation. In fact, this is a major concern in the applied computation literature and has motivated emphasis on Verification. Therefore, Cai et al. (2017) and Cai and Lontzek (2019) developed tools to verify the accuracy of their computations. In some areas of science, there is great concern about the validity of research. To address these concerns, Cai et al. (2017) and Cai and Lontzek (2019) provided Mathematica notebooks that contained the solutions and provide users with both the ability to check the computations as well as use the solutions for their own purposes. This paper will follow both the Verification and Uncertainty Quantification methods and the dissemination practices advocated and implemented in Cai et al. (2017) and Cai and Lontzek (2019). [3]

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. The major shocks in the past 20 years have been the War on Terror, the Great Recession and the fiscal policy responses to the Covid pandemic. We also saw the expansion of Medicare benefits to drugs and a couple of tax cuts. These developments have increased deficits and pushed the level of debt to levels not seen since the end of WWII.

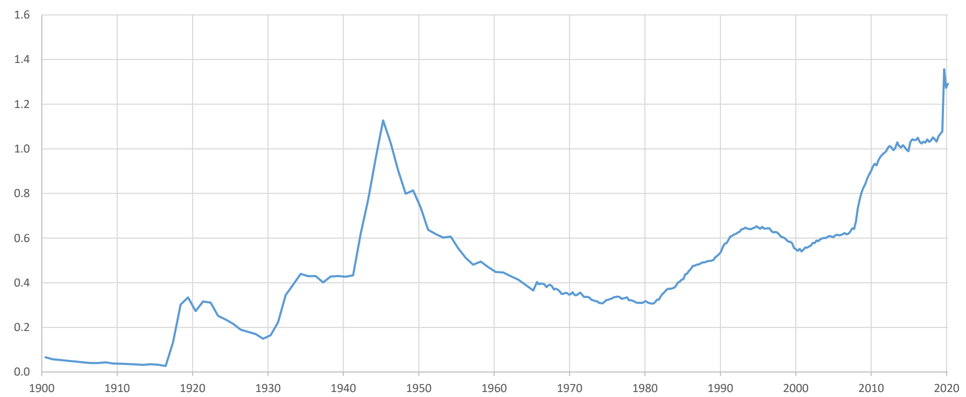

Figure 1: Total public debt as percent of gross domestic product.

[3]This draft currently falls short of meeting those standards. Our focus has been on developing sound and efficient computational methods. Future versions will examine a far broader range of parameter values, results of verification tests, and free computer software which readers can run on their computers to examine the nature and validity of our solutions.

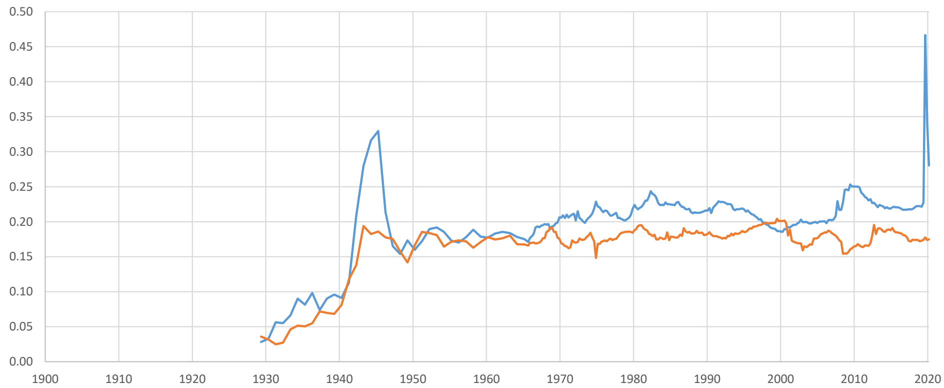

Figure 2: Federal expenditures (blue) and receipts (orange) as percent of gross domestic product.

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.

We depart from much of the macroeconomic literature by assuming that spending is flexible, not fixed by some exogenous stochastic process. It is a common view that the two World Wars of the 20th century were examples of exogenous shocks that forced massive government expenditures. We maintain that government spending is chosen, and need to examine history to justify our assumptions.

The US and UK are the two countries that best fit the assumptions of our models. We contend that any examination of US and UK political history contradicts the assumption of exogenous spending, even when it comes to wars. WWI was not mandatory for either the United States or the United Kingdom. Germany did not attack the UK in 1914. Some argue that the UK was compelled by the Treaty of London 1839 to intervene when German armies moved through Belgium to attack France. However, in "The Pity of War", Niall Ferguson uses documents of UK government discussions to argue that Great Britain knew otherwise. The UK decided to follow its historical policy of blocking any European country becoming the dominant power. In the 18th century, it was John Churchill who lead British forces to block the expansion of French power, and in 1914 it was his descendant, Winston Churchill, who joined those who wanted to block German expansion.

Germany and its allies presented no threat to the US. The US was neutral at the beginning of WWI, although it did stop food shipments to Germany due to the British embargo. These facts plus the special status of Belgium, allowed an American, Herbert Hoover, to organize the Commission for Relief in Belgium, convince Germany of the special status of Belgium and and convince the UK to allow the shipment of food to Belgium. The US tried to be a peacemaker. In his recent book "The Road Less Traveled: The Secret Battle to End the Great War, 1916-1917", Philip Zelikow argues that both sides were exhausted and US diplomatic efforts came close to ending the conflict. The failure of that effort lead to expansion of the war, such as the intensification of German attacks on shipping to the UK and the US entry into the war in 1917. American participation was opposed by many, most notably Senator Robert M. Lafollette of Wisconsin. They argued 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. 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, there was no great desire to enter WWII even after the fall of France. In particular, we need to remember FDR's promise on October 30, 1940, that "Your boys are not going to be sent into any foreign wars". [4]

After the Presidential election of 1940, the US began preparations for entry into WWII and fully embrace Churchill's objective to end German domination of Europe and began planning for the US entry into the European war. Part of this was done in secret. In March, 1941, FDR approved ABC-1, the report of a secret meeting involving American, British, and Canadian military staff that outlined the plan for the European theatre. ABC-1 reads like a history of WWII in Europe.

[4] In fairness to FDR, we note that other speeches that fall included the qualifier "except in case of attack." See http://www.fdrlibrary.marist.edu/daybyday/event/november-1944-9/ .

More publicly, the Lend-Lease act was passed in March, 1941, to help the UK's war effort.

The German invasion of the USSR in June, 1941, changed the strategic situation. Some argued that it weakened the case for US involvement. For example, former President Herbert Hoover argued that the their war would weaken both Hitler and Stalin, and eliminate any serious German threat to either the UK or US. [5] Hoover further argued that Lend Lease should not be extended to the USSR and aid Stalin.

By Autumn 1941, the US and UK had essentially agreed to an alliance to defeat Germany in Europe. Extending Lend Lease to the USSR in November, 1941, made it another ally. The only thing missing was an event that would justify committing American combat forces. Pearl Harbor became the key event. [6]. Hitler's declaration of war against the US on Dec 11, 1941, put the US in the European theatre.

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 should recognize that government spending is chosen, not forced by exogenous circumstances.

3 The Barro (1979) model

Barro (1979) develops a positive theory of government finance. In his model, the government faces a certain, exogenous government spending stream, $\{g_t\}_{t=0}^{\infty}$, that

[5]The events of 1942 help us make that argument. The first American combat with Germany was in North Africa, arguably necessitated by the German threat to British forces in the Middle East and the Suez Canal. German submarine attacks on Atlantic shipping also had to be dealt with. The fact is that those problems were largely solved by mid-1943. The major cost of WWII was the effort to invade Europe and eliminate Hitler.

[6]There was already some combat between US naval vessels and German submarines in October, 1941, but nothing sufficiently serious to trigger US entry. On October 17, 1941, the USS Kearney dropped depth charges on German U-boats while defending a British convoy. The 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.

it finances with taxation and the sale of debt. The government can raise tax revenue, T_t. Public debt takes the form of a one-period, single-coupon bond, b_t, issued at par and outstanding at the end of period $t-1$.

The government incurs the tax collection cost, Z_t, for administration, enforcement, efficiency costs, etc. The collection costs depend positively on the tax collections T_t. The infinitely-lived government chooses taxes to minimize the present value of collection costs. This is a dynamic programming problem with the states including debt level and the state of spending. A quadratic loss function for the government implies that taxes follow a random walk. In other work, Dehejia and Rowe (1995) and Judd (2022) show that the presence of a Laffer curve implies very different dynamics that push government assets (negative debt) up to a level where interest income covers all expenditures. The MSS model surely has a Laffer curve with a finite maximum rate of revenues. We now move to the MSS model, both the fixed expenditure case in their paper and the flexible expenditure level case we focus on.

4 The MSS Model of Optimal Taxation with Flexible or Fixed Spending

The MSS 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 t and provide labor ℓ_t to produce consumption goods c_t and public goods g_t. The supply of ℓ_t is limited to the closed interval $[0, 1]$ and time not spent in formal labor activities, $1 - \ell_t$, is spent at home dedicated to leisure activities or home production. 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_t on public goods. Households spend their income on consumption c_t and one-period bonds b_t.

4.1 Households

The utility of a representative household is a function, $u(c_t, \ell_t, g_t, z_t)$, of consumption c_t, labor supply ℓ_t, and government expenditures g_t. The taste shock z_t follows a finite Markov process with a transition matrix $\pi(z_{t+1}|z_t)$ and affects only how g affects utility. The MSS model assumes that there is a preferred level of government spending, $\bar{g}_(z)$, that depends on the state z. We also assume that u is increasing in consumption, $\frac{\partial u}{\partial c_t} > 0$, and decreasing in time spent working, $\frac{\partial u}{\partial \ell_t} < 0$. The AMSS (2002) paper focuses on the special case where the penalty for deviating from the state-specific target for g is infinite. We will also look at cases where the penalty is quadratic in the deviation. We refer to the MSS specification as the fixed-g case and our specification as the flexible-g case. Therefore, we are formulating the MSS model but generalizing it for flexible g. Of course, we could assume any conventional utility function but the quadratic penalty case is useful for our purposes in this paper.

The consumption, labor, and government expenditure stream, $\{c_t, \ell_t, g_t\}_{t=0}^{\infty}$, is a stochastic process measurable with respect to z_t. This means that z_t is known at the beginning of period t after which $\{c_t, \ell_t, g_t\}$ is determined. The households value the consumption, labor, and government expenditure stream according to 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, risk-free one-period government bonds with a net supply of zero. A bond at time t promises to deliver one unit of consumption at $t + 1$ at the 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} is 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 z_t-contingent 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 reads

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

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,

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

given fiscal policy Φ_t and taste shock z_t, subject to the sequence of intertemporal budget constraints

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

In the following, we assume that the non-negativity conditions do not bind for expositional purposes. The solutions in our examples will always be interior, making the first-order conditions sufficient for optimality. Note that this assumption is not required for our solution approach.

We derive the Lagrangian as

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

where λ_t denotes the Lagrange multiplier on the period-t budget constraint. The household's first-order conditions read

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

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

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

$$\frac{\partial L}{\partial b_{t+1}} : \beta E_t[\lambda_{t+1}] - p_{b,t}\lambda_t = 0. \tag{7}$$

This is a standard stochastic control problem representing the households' reactions to government policy When the government determines its policies, it must consider the effect of the households' expectations of its future policies.

4.2 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

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

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 expenditures 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 τ_t, transfer tr_t, and spending decisions g_t and chooses the market price for bonds $p_{b,t}$, the shadow price of consumption for the next period, and the recommended consumption, labor and bond choices for the household. Then, the households solve their optimization problem subject to the government's fiscal policy choice and pick the allocation suggested by the government if it is in their interest. To ensure that households follow through with the plan, the government chooses fiscal policy and household allocations consistent with the households' incentives.

The government's problem can be written as

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

subject to its budget constraint

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

its aggregate resource constraint

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

the first-order conditions from the household's problem

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

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

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

$$\beta E_t \lambda_{t+1} - p_{b,t} \lambda_t = 0. \tag{14}$$

Note that we still ignore the non-negativity constraints for expositional purposes.

5 Problem Formulations

In this section, we first present the first-order approach taken by MSS and briefly point out the shortcomings of the MSS approach. We then present a dynamic programming formulation.

5.1 First-Order Approach

The first-order approach starts from the government's optimization problem and takes the approach used in Turnovsky and Brock (1980). 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 λ—by IC^c, IC^ℓ, IC^b, and IC^λ, respectively. For notational brevity, $U_{t,x}$ denotes the partial derivative of U_t with respect to x_t; analogously, $U_{t,xy}$ denotes the second order partial derivative of U_t with respect to x_t and y_t.

The Lagrangian of government problem is

$$L = E_0 \left[\sum_{t=0}^{\infty} \beta^t \Big[U_t - \theta_t^{GB} GB_t - \theta_t^{PPF} PPF_t - \varphi_t^c IC_t^c - \varphi_t^\ell IC_t^\ell - \varphi_t^b IC_t^b - \varphi_t^\lambda IC_t^{\lambda_t} \Big] \right]; \tag{15}$$

the first-order conditions are

$$\frac{\partial L}{\partial c_t}: \quad U_{t,c} - \theta_t^{PPF} - \varphi_t^c U_{t,cc} - \varphi_t^\ell U_{t,\ell c} - \varphi^\lambda = 0 \tag{16}$$

$$\frac{\partial L}{\partial \ell}: \quad U_{\ell_t} - \varphi_t^{GB} \tau_t - \varphi_t^{RC} - \varphi_t^c U_{t,c\ell} - \varphi_t^\ell U_{t,\ell\ell} - \varphi_t^\lambda (1 - \tau_t) = 0 \tag{17}$$

$$\frac{\partial L}{\partial b_{t+1}}: - p_{b,t} \theta^{GB} + \beta \theta^{GB} = 0 \tag{18}$$

$$\frac{\partial L}{\partial g_t}: \quad U_{t,g} + \varphi_t^{GB} + \varphi_t^{RC} - \varphi^c U_{t,cg} - \varphi^\ell U_{t,\ell g} = 0 \tag{19}$$

$$\frac{\partial L}{\partial \tau_t}: \quad - \varphi_t^{GB} \ell + \varphi_t^\ell \lambda_t = 0 \tag{20}$$

$$\frac{\partial L}{\partial \lambda_t}: \quad \varphi^c - \varphi^\ell (1 - \tau_t) + \varphi^{EE} p_{b,t} = 0. \tag{21}$$

The appearance of the cross-derivatives like $U_{t,cl}$ implies that this optimization problem may not be convex. The failure of convexity necessitates global optimization methods.

5.2 Recursive Formulation

We now present a formulation of the dynamic programming approach by Kydland and Prescott (1980) that applies to economic planning problems. This approach corrects the claim by Prescott (1975) that optimal control methods cannot be applied to economic planning. We use the dynamic programming formulation of the recursive problem described in Kydland and Prescott (1980). Their approach treats the current bond level b_t, the shadow price λ_t, and the current shock z_t as state variables. At any time, these are sufficient statistics for the dynamic optimization problem. In the following, we drop the t subscript for notational brevity and denote the next period's variables by a $+$ superscript.

Each period the government chooses c, l, b^+, g, τ, p_b and the next period's values for $\lambda^+[z^+]$ Since λ is measurable with respect to the z process, a different value for λ^+ can be chosen for each potential value of z^+. This is all implied by the optimization problem specified by the first-order approach. Using the language of principal-agent theory, the government promises the next period's values for λ^+ conditional on z^+. To clarify this connection, we say that the government chooses $\lambda^+(z^+)$ for each possible z^+.

At this point, we emphasize that the first-order approach and this dynamic programming approach are solving the same problem. They both assume precommitment by the government. This is clear in the first-order approach because, at time 0, we choose the time series for each economic variable. The solution to our dynamic programming problem will also be the precommitment solution because it uses the same state variables tracked by the solution to the Lagrangian problem. The solution for this problem may be time-inconsistent. We do not study the differences between the optimal solution with precommitment and a Nash equilibrium set of policies when there is no commitment.

The government's dynamic programming problem is

$$\max_{c,l,g,\tau,b^+,\lambda^+(z+)} \quad u(c,l,g,z) + \beta \mathbb{E}[V(b^+, \lambda^+(z^+), z^+] \tag{22}$$

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

$$-\lambda + \frac{\partial}{\partial c} u(c,\ell,g,z) = 0 \tag{23}$$

$$(1-\tau)\lambda + \frac{\partial}{\partial \ell} u(c,\ell,g,z) = 0 \tag{24}$$

$$-b + c - tr - \ell(1-\tau) + p_b b^+ = 0 \tag{25}$$

$$\beta \sum_{z^+} \lambda^+(z^+)\pi(z^+|z) - p_b\lambda = 0 \tag{26}$$

$$\tag{27}$$

Not all state variables $\{b, \lambda\}$ allow for a policy sequence that fulfills the household's incentive compatibility, the government budget constraints, and the aggregate resource constraint for every possible realization of the government spending taste. We denote the values of (b, λ) which are possible for spending state z by $\Omega(z)$ and refer to it as the feasible region given z. Our dynamic program implicitly assumes that $(b+, \lambda(z+)) \in \Omega(z+)$, and implements this by setting $V[b, \lambda, z] = -\infty$ for any (b, λ) not in $\Omega(z+)$. One step in Section 5.3, is to approximate $\Omega(z)$.

MSS uses PEA to solve the model. We have examined their code and found that it often does not converge. Also, the solution depends on the seed they use for the random number generator. When it does converge, the long run debt level appears to become very negative, implying a war chest in the long run. They also never check the first-order conditions of their solution. Instead, they choose a solution which minimizes the sum of squared errors. Our dynamic programming approach will specify that all first-order conditions hold at all times.

5.3 Computational Algorithm

Solving the stochastic dynamic policy problem of the government presents many computational challenges. First, the maximization problem may not be convex and, second, the feasible regions $\Omega(z)$ are unknown, endogenous, and may be non-convex. The value function will diverge to $-\infty$ for feasible states close to the boundary of the feasible region, substantially complicating the approximation of the value function. We address these issues in this section.

5.3.1 Discretization

We discretize the state space in order to avoid the difficulties of approximating the value function. We use a fine grid of points and verify the quality of the discrete approximation. Discretization also allows us to find global maxima by exhaustive search. 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 shadow price λ uniformly as well. However, in the case of additively separable utility functions that satisfy the Inada conditions, λ tends towards infinity for $c \to 0$. Therefore, instead of discretizing λ uniformly, we discretize the consumption c uniformly starting from some lower bound c_{min} up to $c_{max} = 1$. This implies a non-uniform discretization of λ. Discretizing c means that both state variables are directly observable variables with the same units, a feature which makes it easier to interpret the results. We do not treat the case of nonseparable utility in this paper but believe that a similar transformation could be found for those cases.

5.3.2 Feasible Region

To determine the numerically feasible region $\tilde{\Omega}(z)$, a subset of the feasible region $\Omega(z)$, we iteratively improve the k-th approximation of the feasible region $\tilde{\Omega}^k(z)$

until we find its fixed point, i.e., $\tilde{\Omega}^{k+1}(z) = \tilde{\Omega}^k(z)$.[7] The implementation of this process is straightforward: As a starting guess, we choose a "stay where you are" policy function, that is, $b^+ = b$ and $\lambda^+(z^+) = \lambda$. This yields an initial guess for the feasible region denoted by $\tilde{\Omega}^0(z)$. In iteration k, we improve upon the previous guess $\tilde{\Omega}^{k-1}(z)$ by searching for a feasible policy for every infeasible state, producing $\tilde{\Omega}^k(z)$. It is sufficient to search in the set of states declared feasible in the step $k-1$, that is, $\Omega^k(z) \setminus \Omega^{k-1}(z)$ for $k > 1$, $\Omega^0(z)$ otherwise. This iteration continues until we have reached the fixed point $\tilde{\Omega}^k(z) = \tilde{\Omega}^{k+1}(z)$ for all z. [8]

5.3.3 Algorithm and Implementation

In this subsection, we briefly outline the steps of our computational framework. Once we have determined the feasible region, we solve the dynamic programming problem. We use the fact that we can treat the dynamic programming problem in Equation (22) as a bi-level optimization problem. First, we fix the dynamic choices (b^+, λ^+) and solve the resulting static problem:

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

[7]Note that the numerically feasible region $\tilde{\Omega}(z)$ is a subset of the actual feasible region $\Omega(z)$ due to the approximation of the state space. Our results indicate that this approximation error is of no importance because the optimal policy pushes the state away from the boundary.

[8]We have also developed a complementary procedure that begins with some clearly infeasible states and then finds other clearly infeasible states. The two procedures produce the same feasible set for the examples we solve.

where $U(b^+, \lambda^+)$ is the maximized utility from the conditional on a fixed b^+ and λ^+. In some cases, this inner, static problem can be solved in closed-form. The outer level, dynamic problem reads

$$V(b, \lambda, g) = \max_{b^+, \lambda^+} \; U(b^+, \lambda^+) + \beta \mathbb{E}[V(b^+, \lambda^+, g^+)] \tag{28}$$
$$s.t. \;\; (b^+, \lambda^+) \in \Omega(z^+),$$

We proceed in a standard way for discrete-state dynamic programming. We first guess a value function. In the policy improvement step, we globally search for the state transition $(b^+, \lambda^+(z+)) \in \Omega(z^+)$ that maximizes the outer, dynamic problem (28) for each state (b, λ). For the case of separable utility, we can compute what is essentially the closed-form solution of the inner, static problem. More precisely, we precompute an analytic, highly accurate, approximation of the solution. This allows us to implement the global search on graphical processing units (GPUs). With the new policy, we compute the value function which arises if we follow the new policy forever. We continue this policy iteration method until convergence.

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 used in our results below was not significantly different from coarser discretizations regarding 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.

6 Parameterization

We follow the literature by assuming an additively separable utility function

$$u(c,\ell,g,z) = \underbrace{\frac{(c+\underline{c})^{1-\sigma_1}}{1-\sigma_1}}_{\equiv uc(c)} + \eta \underbrace{\frac{(1-\ell)^{1-\sigma_2}}{1-\sigma_2}}_{\equiv ul(l)} - \underbrace{\theta(g-g_z)^{\sigma_3}}_{\equiv ug(g,z)}, \tag{29}$$

where $\underline{c}$ denotes a fixed level of consumption every household consumes. The parameters σ_1, σ_2 and σ_3 are utility function parameters; η and θ are multipliers of the disutility of labor and deviations of the government spending from its taste, respectively. We assume that total time for each agent is 1, which implies that the total possible GDP is 1. All quantities are expressed as fractions of total possible GDP. For example, $g = 0.1$ means that government expenditures is a tenth of total possible GDP.

The additively separable utility function specification (29) allows for closed-form solutions of the consumer's problem. The optimal consumption equals the inverse marginal utility of consumption of the shadow price λ,

$$c_t = \frac{\partial}{\partial c} uc^{(-1)}(\lambda), \tag{30}$$

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

$$\ell_t = \frac{\partial}{\partial l} ul^{(-1)}(-(1-\tau)\lambda). \tag{31}$$

The price of tomorrow's bond the household is willing to pay depends on the expected discounted shadow price tomorrow λ^+ normalized by the shadow price today λ as

$$p_b = \beta \frac{\sum_{z^+} \lambda^+(z^+)\pi(z^+|z)}{\lambda}. \tag{32}$$

Note that this parameterization allows for a direct link between the consumption c and the shadow price λ. Since consumption is the more intuitively understandable variable, we use consumption for most of the plots instead of λ.

In this paper, we assume $\sigma_1 = 1$, $\sigma_2 = 1$, and $\sigma_3 = 2$. The appendix covers

Table 1: Utility parameters

σ_1	σ_2	σ_3
0.75	0.5	2
0.75	1	2
1	0.5	2
1	1	2
1	2	2
2	0.5	2
2	1	2
2	2	2

some other cases (see table 1). Equations (30-32) neglect the non-negativity constraints on c and ℓ for expositional purposes.

We set $\beta = 0.96$ to facilitate comparisons with MSS. We assume $\underline{c} \in \{0, 0.1\}$, whereas $\underline{c} = 0$ corresponds to model proposed by MSS. Assuming $\underline{c} > 0$ allows the (market) consumption c to equal 0 in the presence of high taxes due to a finite marginal utility of consumption for $c = 0$, implying that revenue maximization is attained at a tax rate below 1; that is, it produces a realistic Laffer curve. If $\underline{c} = 0$ then for many utility functions the tax rate can be close to 1 with agents consuming little leisure or consumption. This is not plausible for any economic system we want to consider.

To facilitate comparisons between the fixed-g and flexible-g cases, we assume that g is chosen according to a utility function that penalizes any deviation from a target level of spending where the target follows an exogenous Markov process. We set $\theta = 100$, which represents a situation that highly penalizes deviations from the government spending target. Changes in this target is a taste shock, but the flexibility in government spending allows the government to reduce its spending in (unexpectedly) long-lasting periods where the desired level of spending is high. In some way, this specification implies exogenous government spending if the penalty for deviating from te target is infinity. However, we find that the exogenous spending case is not the limit as the penalty for deviations becomes infinite. Exogenous spending is a fundamentally different, and unrealistic, assumption.

We allow two values for g_z, the target expenditure, which we call peace and war. We follow Buera and Nicolini (2004) in choosing the following transition probabilities between peace and war:

$$\Pi = \begin{bmatrix} 0.9787 & 0.0213 \\ 0.3333 & 0.6667 \end{bmatrix}. \tag{33}$$

This roughly matches the US experiences in the last century: approximately two major wars per century with an average duration of 3 years.

7 Feasible Regions for the Log-Log case

The feasible level of government debt is central to the debate on public debt levels. This section contrasts the feasible regions for the fixed-g with the feasible regions under our flexible-g specification. Our results imply that they differ fundamentally: the fixed-g parameterization implies almost no feasible government debt levels, whereas the flexible-g parameterization allows for significant levels of government debt.

The feasible region depends on $\underline{c}$. Setting the fixed level of consumption $\underline{c}$ to 0 (as done in MSS) drastically extends the feasible region and allows high levels of government debt. This implies an infinite marginal utility of consumption at zero consumption, making it impossible for agents leave the market no matter how high taxes go. We want to examine cases where there is an empirically reasonable maximum level of revenue. Choosing positive values for $\underline{c}$ makes that possible. We chose $\underline{c} = 0.1$

7.1 Endogenous Government Spending

We next describe the feasible region of the flex-g case. The key fact is that zero spending is always possible, implying that the feasible region does not depend on the shocks to tastes for g. This is illustrated in Figure 3. We also see that significant debt levels are possible. In our example, the average GDP is about 1/3 implying that a debt level of 5 is 15 times GDP. This example is meant to illustrate key ideas, not to be quantitatively plausible.

Figure 3 displays three regions. The dark brown region display states that are war-chest states in peace and war. That is, the revenue from government assets covers all spending with the excess revenues rebated lump-sum to agents. The light brown region are states that are war-chest states in peace, but not in war. The blue states are all feasible states but distortionary taxes must be levied. The government can always choose to spend less even though this is heavily penalized. Therefore, the feasible space includes all levels of debt that can be financed while zeroing out any spending.

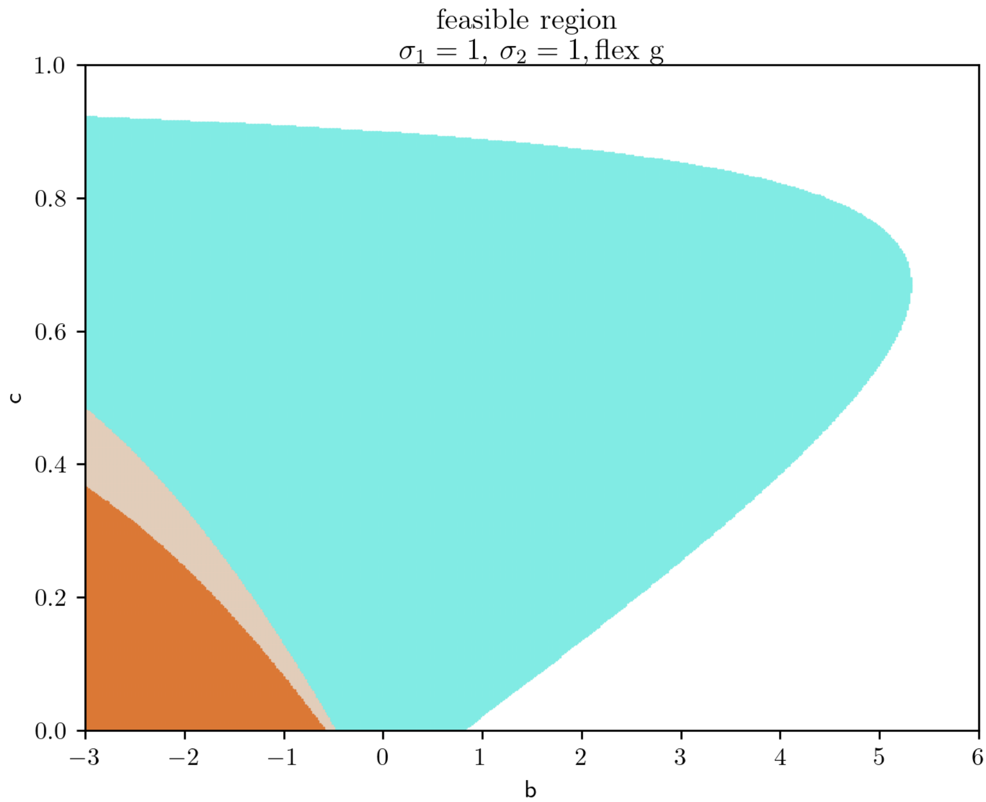

Figure 3: Feasible states with flexible g

7.2 Exogenous Government Spending

Figure 4 and figure 5 displays the feasible region for two different exogenous government spending processes. We first fix government spending in peace to 0.09

and government spending in war to 0.27. Consumption is typically around 0.30 implying historically reasonable ratios between consumption and g. In the high spending case—where peace spending equals 0.09 and war spending 0.27—almost no government debt is feasible. The broken line is the boundary of the feasible set for flexible g.

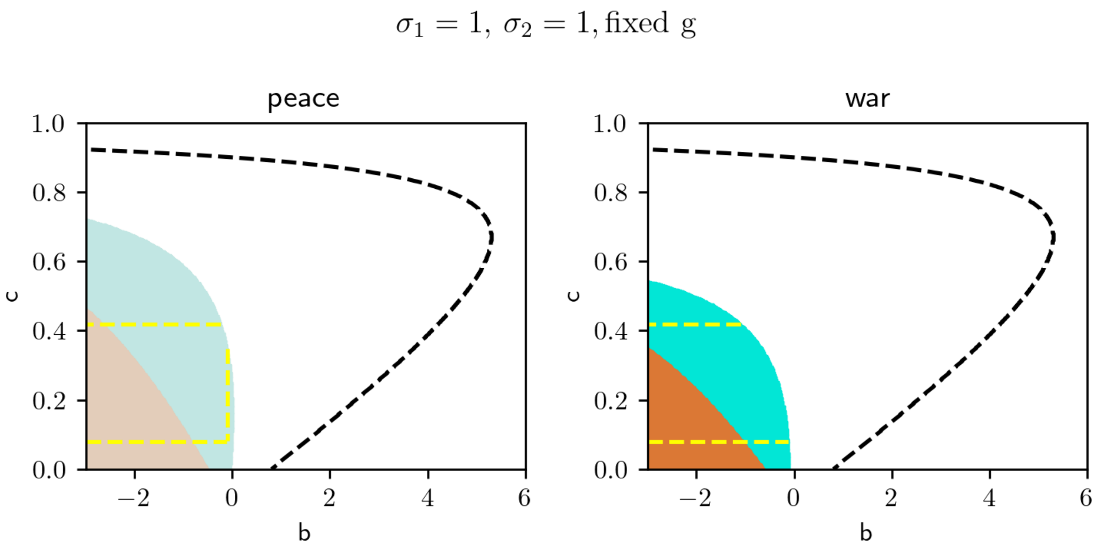

Figure 4: Feasible states with fixed g

Our next plot shows that positive levels of debt are feasible if we reduce the level of g by half. As one might expect, the feasible region increases for lower levels of war spending.

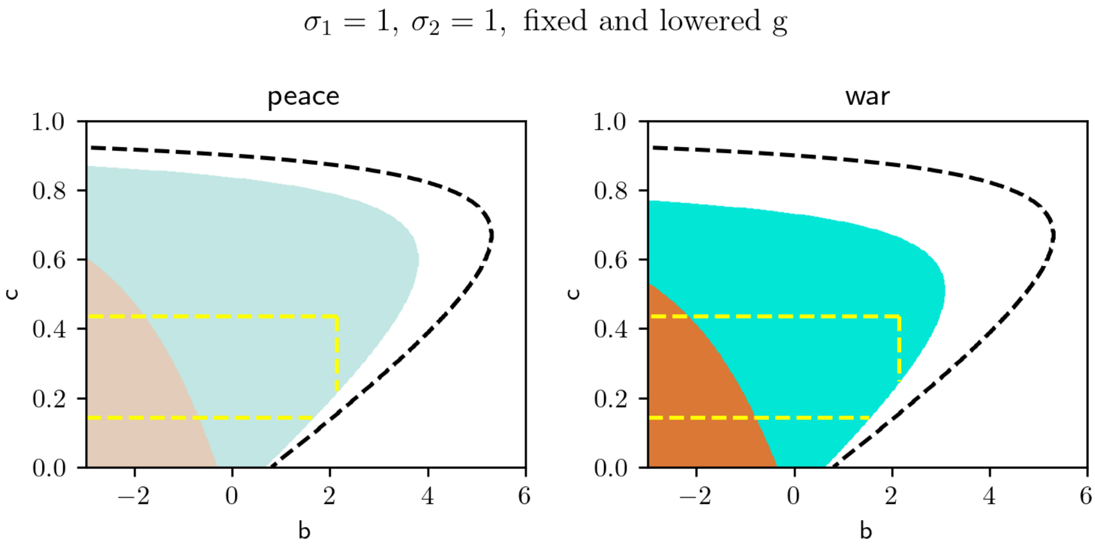

Figure 5: Feasible states with less fixed g

The government can only be in debt in the case of the low government spending process.

8 Value function and Excess Burden of Taxation

Figure 6 displays the value function in terms of the (b, c) states. It is defined only over the set of feasible (b, c) states. Note that the value function drops rapidly as one approaches the boundary of the feasible region.

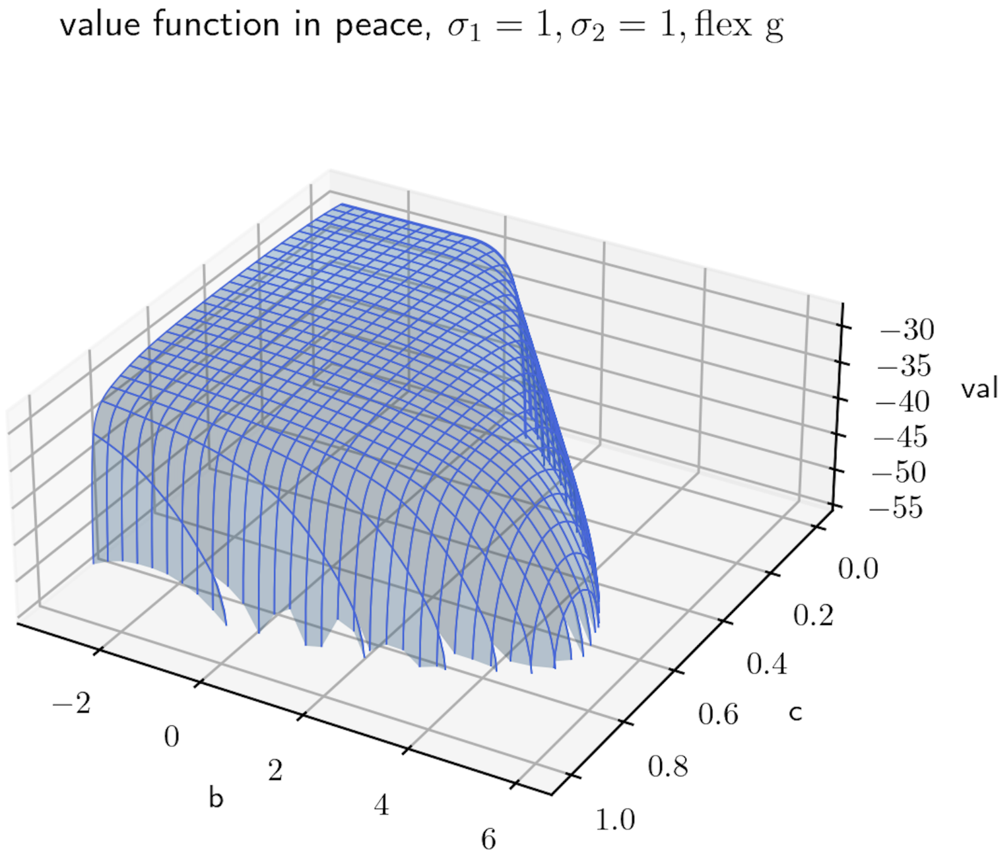

Figure 6: Value function

All the numbers in our plots and graphs will be rather artificial. The next plots relates our stylized model to general properties of taxation. Figure 7 and figure 8 and are contour plots depicting the excess burden of raising one dollar of tax revenue in peace and war. If taxes are lump sum, the excess burden is zero. The excess burden measures the welfare cost of marginal distortionary taxation. If debt is low, then taxes are low and the excess burden is small. In our model, b=1 corresponds to a debt level 3 times GDP, but the excess burden is about 1 or higher. The estimates of excess burden of real world taxation vary greatly depending on what tax we change, but excess burdens on the order of -0.5 to -2.0 are common. The excess burden is a measure of the tradeoff a government has between taxation and spending. Therefore, regions where b is between 1 and 2 are sensible for us to examine when comparing our model to real economies.

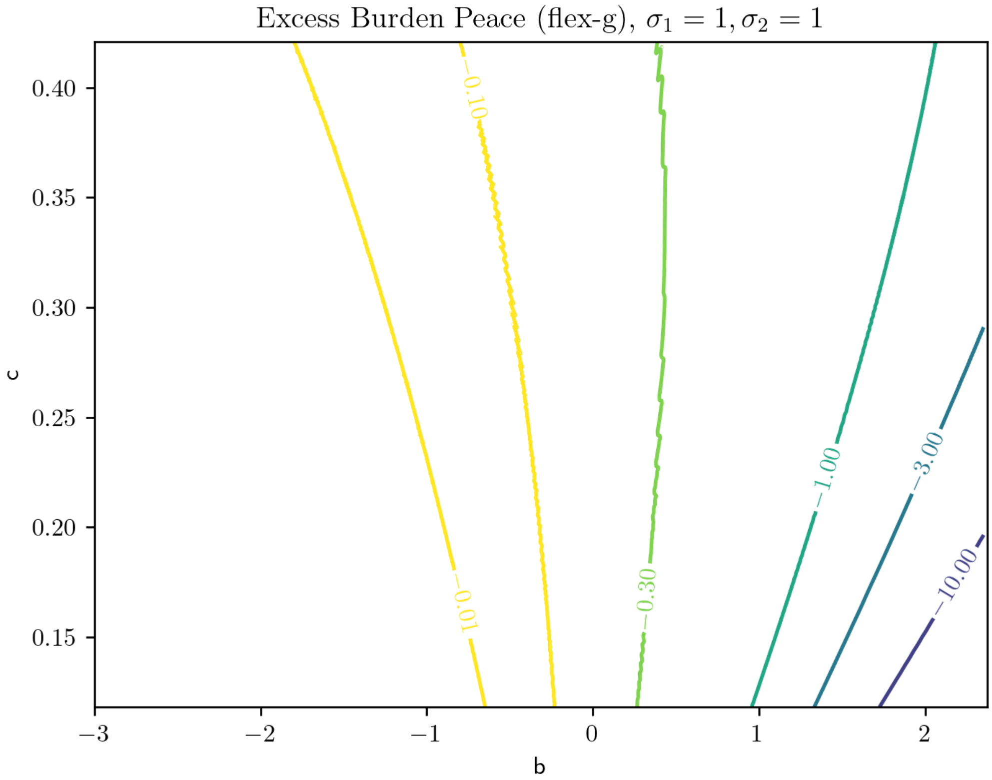

Figure 7: Excess burden contours in peace

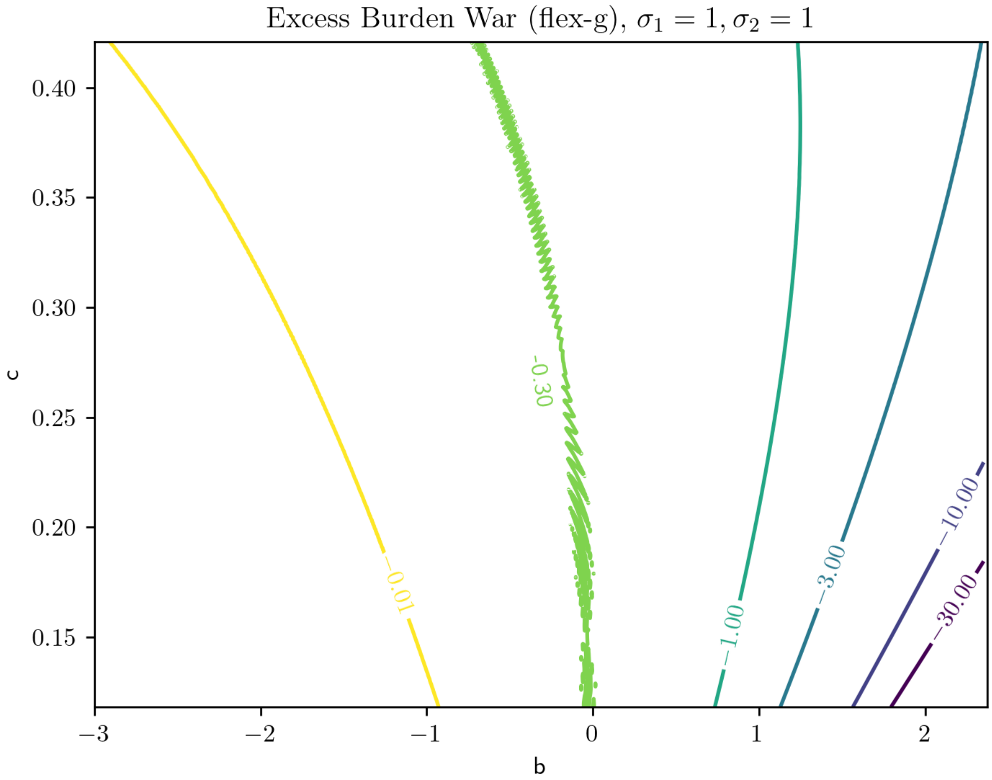

Figure 8: Excess burden contours in war

9 Log-Log case

In the paper we focus on one specific utility function where $\sigma_1 = \sigma_2 = 1$. We first examine dynamics of the flex-g model in three scenarios: perpetual peace, perpetual war, and random transition following the transition matrix Π. Each scenario is described by a collection of six subgraphs. We examine the evolution of the state for 12 alternative initial values for (b, c). The left-hand (right-hand) side plots display the evolution of debt (consumption). The three rows correspond to three values for c0, the initial consumption. Those values are c0=0.421, 0.270, 0.118. Each plot considers up to five initial states for debt, b. The values for debt are b0= -2, -0.906, 0.187, 1.281, 2.375. These paths help us understand how the system evolves for the simple case of perpetual peace. In contrast to Barro and MSS, we find no simple description, neither a tendency to build a war chest nor a random walk.

9.1 Perpetual Peace/War simulations

The next plots display paths of the state variables under assuming perpetual peace. We expect these paths to be similar to paths in general because the probability of going from peace to war is 0.02. In all cases, we see that debt is higher when the initial level of debt is higher. Similarly, consumption is higher when c0 is higher. These are not major findings but is rather a test of the soundness of the algorithm. If the state space was too coarse, then one could see some of these paths crossing. Our discretization of b and c is sufficiently fine that we see smooth paths for both states.

There is a tendency for debt to fall, sometimes falling to a negative level of debt such that interest on government assets covers any expenditures and there is no need for taxes. The absence of a tax in the long run is seen in the consumption plots where consumption hits the first best level, which is slightly above 0.4. However, this is not true when debt is initially high and c0 = 0.270. In that case, debt settles in at a constant positive value slightly less than 2, implying significant taxation which then leads to the low level of consumption. Furthermore, we see that scenario 5 is missing when c0=0.421 and when c0=0.118. This because that combination of b and c is not feasible.

Perpetual Peace (flex-g), $\sigma_1 = 1, \sigma_2 = 1$

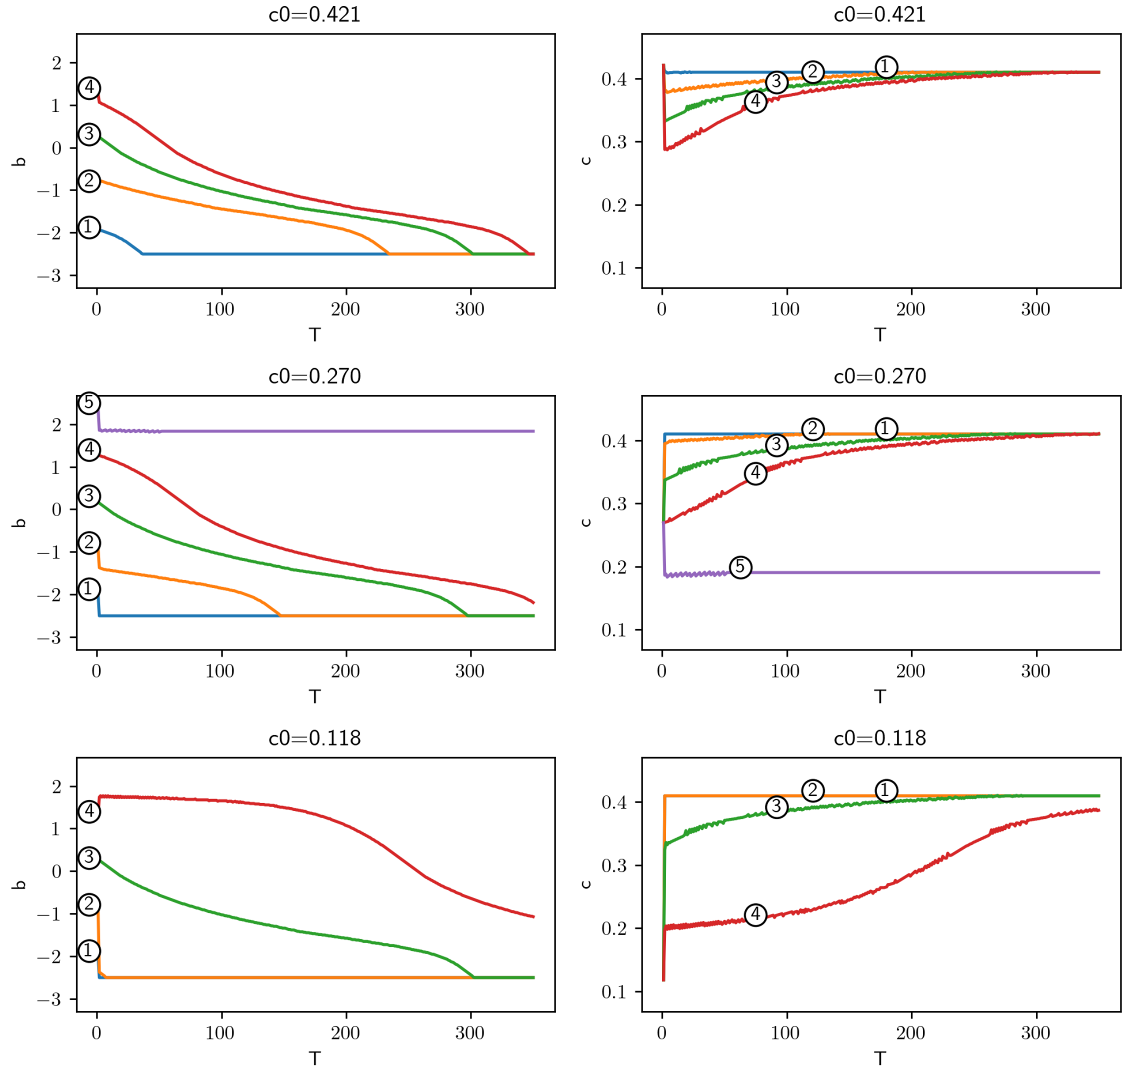

Figure 9: The policy paths and optimal choices for tax rate and government spending over 55 periods of perpetual peace. All simulations start at an initial level of consumption of $0.421, 0.27$ and 0.118 with initial government debt levels of -3.03, -2.06, -0.11, and 1.83, plotted by the blue, red, green, and purple lines, respectively.

The second plot displays the same results if there is perpetual war. The results are again natural with war expenditures forcing debt upwards except for the one case where b0 is -2 and c0=0.118. In that state, taxes are high enough to push the debt level down to a war chest level and have all spending financed by interest on assets after that.

The intuition is clear. When in war, there is a 1/3 chance of peace in the next period. Therefore, it is natural to use debt to finance a level of expenditure that is considered high relative to future expenditures, smoothing the tax rate. However, if the war continues, the rise in debt hits a ceiling because the government does not want debt to become so great that financing it is infeasible. In this situation, the government cuts back on its expenditures. If b starts sufficiently low, the government moves into a war chest region. In this region, no taxes are necessary, and the interest income allows for sustained high government spending. However, if b starts higher, b climbs slowly to a limit, which is roughly the same for all initial levels of b. At the same time, the tax rates increase, and government spending decreases substantially.

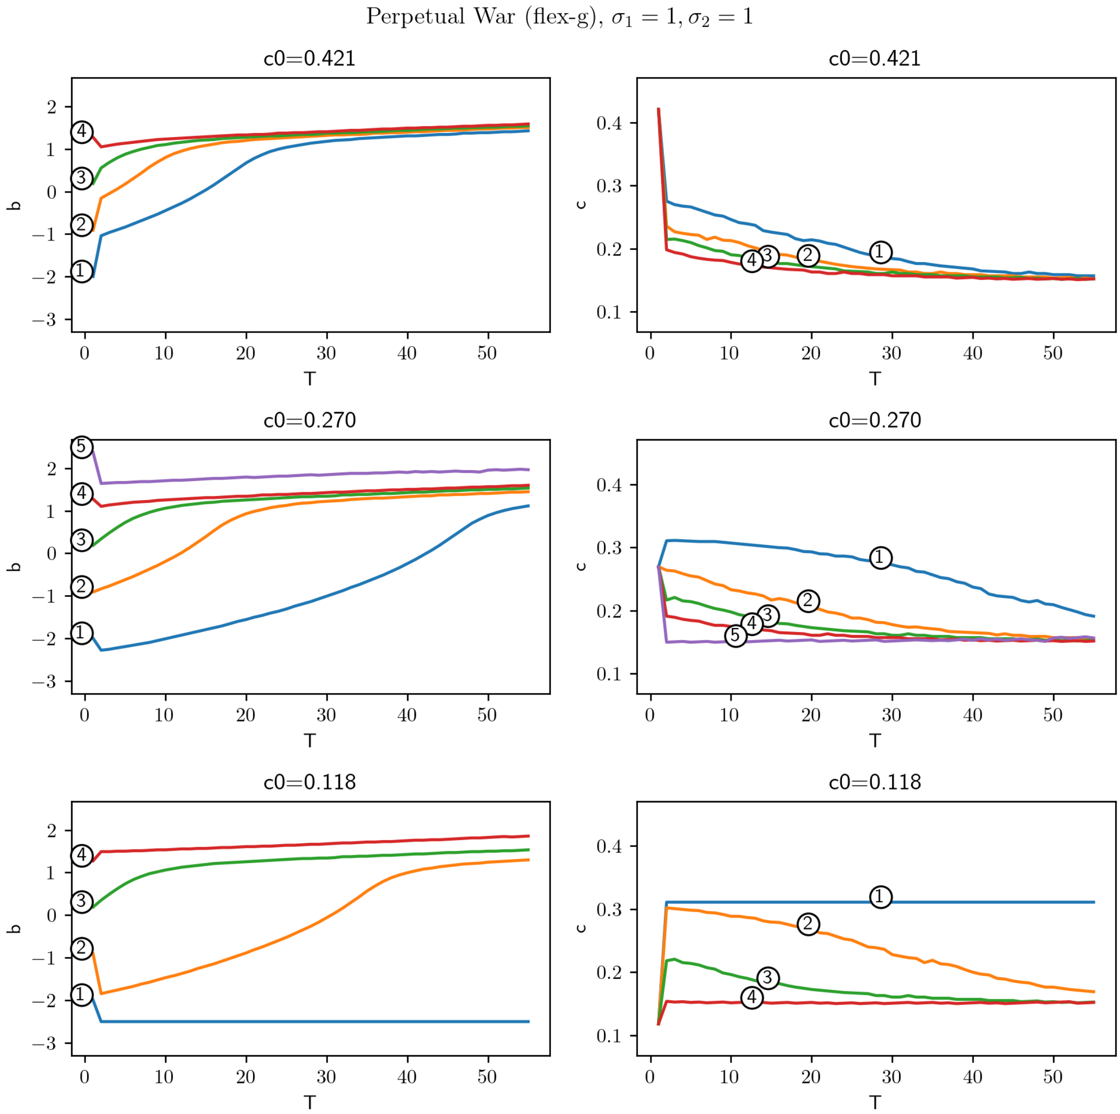

Figure 10: The policy paths and optimal choices for tax rate and government spending over 55 periods of perpetual war. All simulations start at an initial level of consumption of $0.421, 0.27$ and 0.118 with initial government debt levels of -3.03, -2.06, -0.11, and 1.83, plotted by the blue, red, green, and purple lines, respectively.

9.2 Random Paths

Next, we analyze the time series of government debt b, the shadow price λ, the government spending g, and the tax rate τ for a stochastic path of peace and war, generated according to the Markov transition matrix Π.

ALICE:WHERE ARE THE RANDOM PATH PLOTS? In Figure **??**, the government starts with low government debt at the state ($b_{init} = 1.0, \lambda_{init} = 2.4$). After a peak in debt, the government builds up assets of $b = -1.5$, which the government mostly sustains for the entire horizon—–the shadow price λ jumps between 2 and 2.6 for peace and war, respectively. The government approximately meets its spending taste for peace and war. In peace, the tax rate of about 0.05 is relatively low and is raised to 0.35 in war.

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 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.

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

9.3 Transition dynamics

Figure 11 displays the (b, c) state transitions in four situations: current state is peace and next period is peace (called "peace to peace"), peace to war, war to peace, and war to war. Each arrow represents the path from the state in the previous period to the current state. The state transitions display a clear pattern: If c is larger than 0.4 or lower than 0.15, the next period's consumption transitions to states $c \in [0.15, 0.4]$; if c is in $c \in [0.15, 0.4]$, it stays in this interval. Therefore, we refer to the set of all states with consumption $c \in [0.15, 0.4]$ as the ergodic box

$E = \{(b, c) : c \in [0.15, 0.4] \wedge (b, u_c(c)) \in \Omega\}$. The ergodic box is denoted by broken yellow lines. We see that the optimal policy will push the system to points in that box from any initial state outside that box.

The first plot displays the law of motion when we move from peace to peace for various states outside the box with the broken yellow boundary. This true for any peace/war scenario.

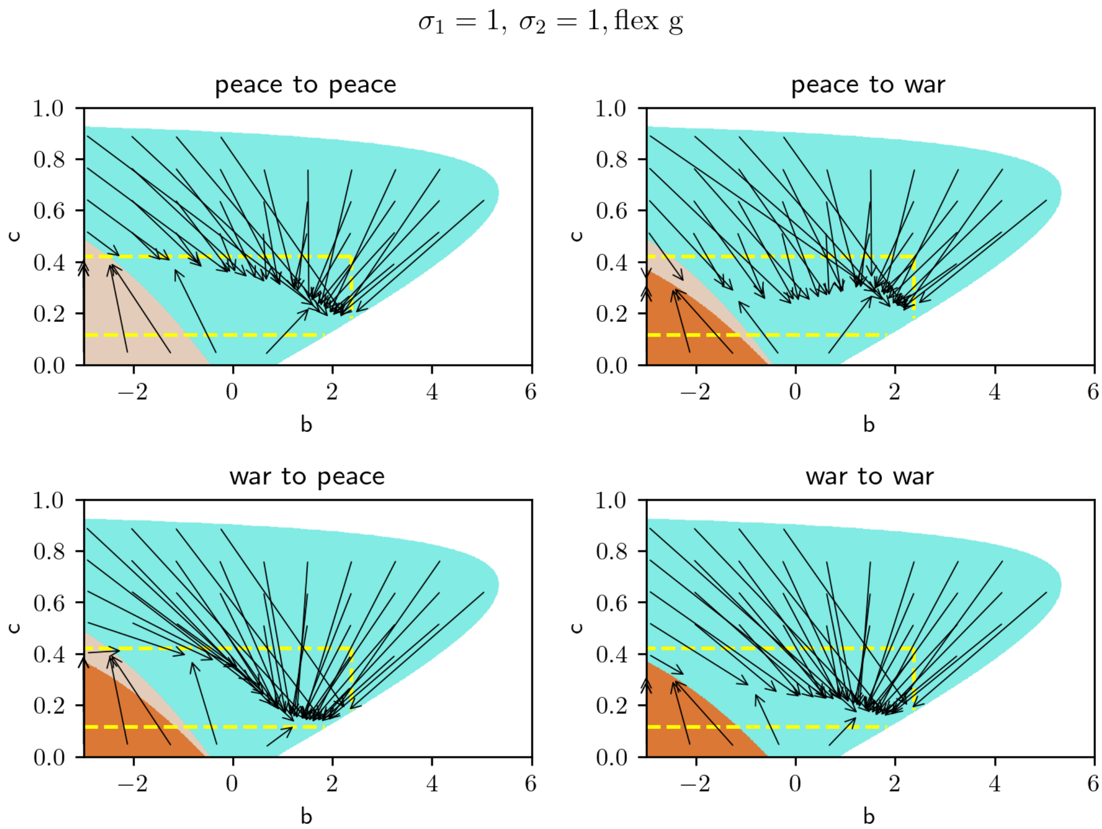

Figure 11: State transitions for points outside of ergodic box.

It is also true that the solution will keep the state inside the ergodic box. Therefore, we will confine our attention to dynamics within the ergodic box. Note that we assume $b > -5$. This has no impact because asset levels that large implies interest covers all expenditures and we assume remaining revenue finances a lump-sum transfer. The next animations display the law of motion for the state variables inside the ergodic box. Click ">" to see slices of the law of motion. That is, each slice has one initial bond level but many possible initial consumption levels. As we move through the various slices, we see that for each peace/war scenario there is a collection of states to which the state moves. The key feature is that any state in the ergodic box moves to a point in a much smaller set of states, denoted by the black dots. We will exploit that structure in the next section.

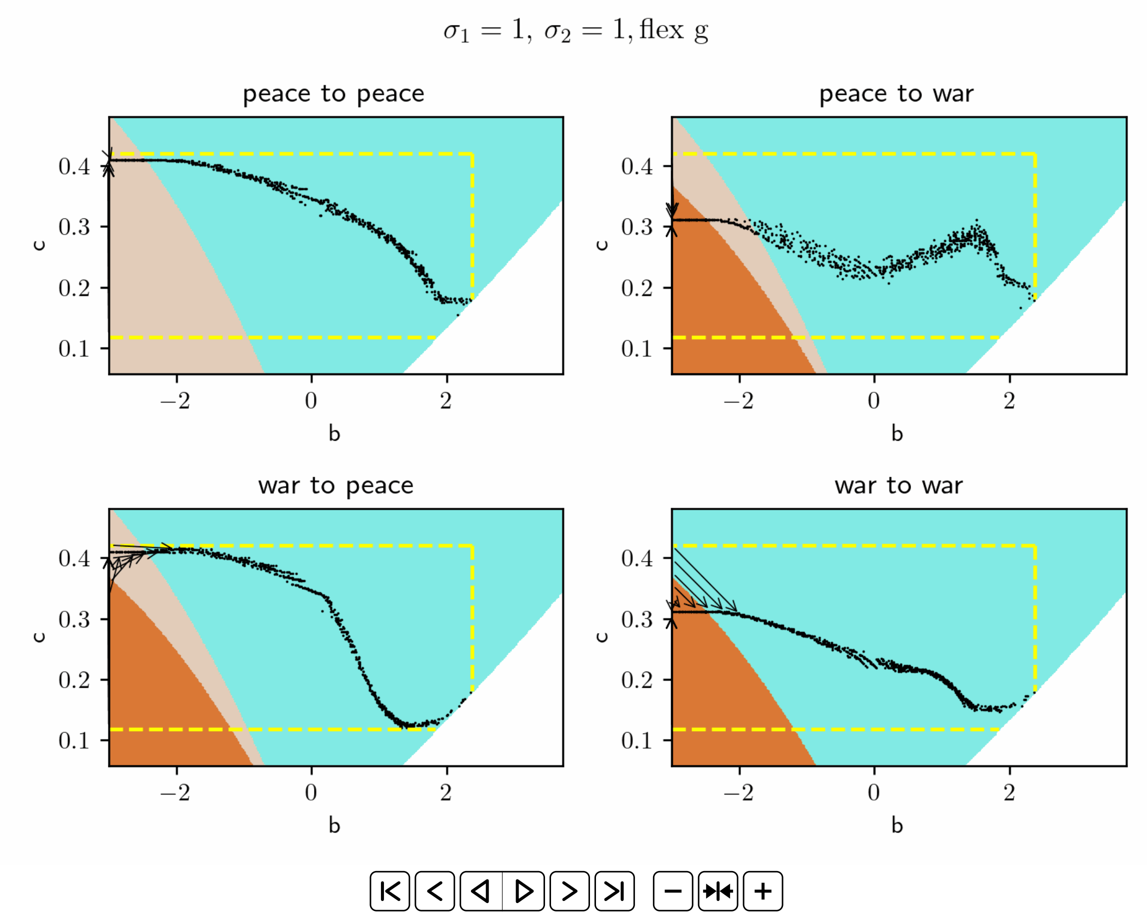

9.4 Ergodic Sets

The optimal tax policy does not follow a stationary process in any usual sense of the term. Since the policy lives on a finite state, the resulting Markov process is stationary in some long run sense but the long-run can be so far in the future that the concept has no value for thinking about the optimal policy. This requires us to use different tools to describe the process for the economic variables that comes from the optimal policy. The first step is to focus on the states that are not transient. We first saw that points outside the ergodic box are transient. We next define what we call the "ergodic set". This may be an abuse of the term because we do not intend to compute the long-run distribution of the state. We instead want to define a set which will surely contain the process after some initial time. Our ergodic set is the fixed set for the transition rule; that is, if the current state is in this set then the next period's state will also be in the set, and each point is reachable from some other point. [9]

The next plot displays the ergodic set. It consists of four kinds of states, each kind represented by a different color. If the current government spending state is peace, then the state is either a green or yellow dot, green if the state was peace in the previous period and yellow if war. The red and purple dots represent states visited in war, red if the previous period was war and purple if peace. Each color represents one of the four different scenarios representing the spending state today and yesterday. If (b,c) is a green dot, then we know that (b,c,peace) is a frequently visited state and that the previous period was also peace, whereas if it is yellow then (b,c,peace) is a frequently visited state but only if the previous period was war. Similarly, red is the color of the states if both today and yesterday were war, and purple is when today is war but yesterday was peace. The collection of all colored dots is the true ergodic set of the process in the sense that if you are at such a state then you will move to another such state.

The small size of the ergodic set implies special features of the dynamics. Even though the state space has two dimensions, the ergodic set is composed of four

[9] There may be multiple such fixed sets. We focus on the largest fixed set.

distinct sets, each of which is close to having just one dimension. In particular, we see that debt and consumption are strongly correlated conditional on yesterday's and today's peace/war state. We will exploit that below when we describe how policies evolve.

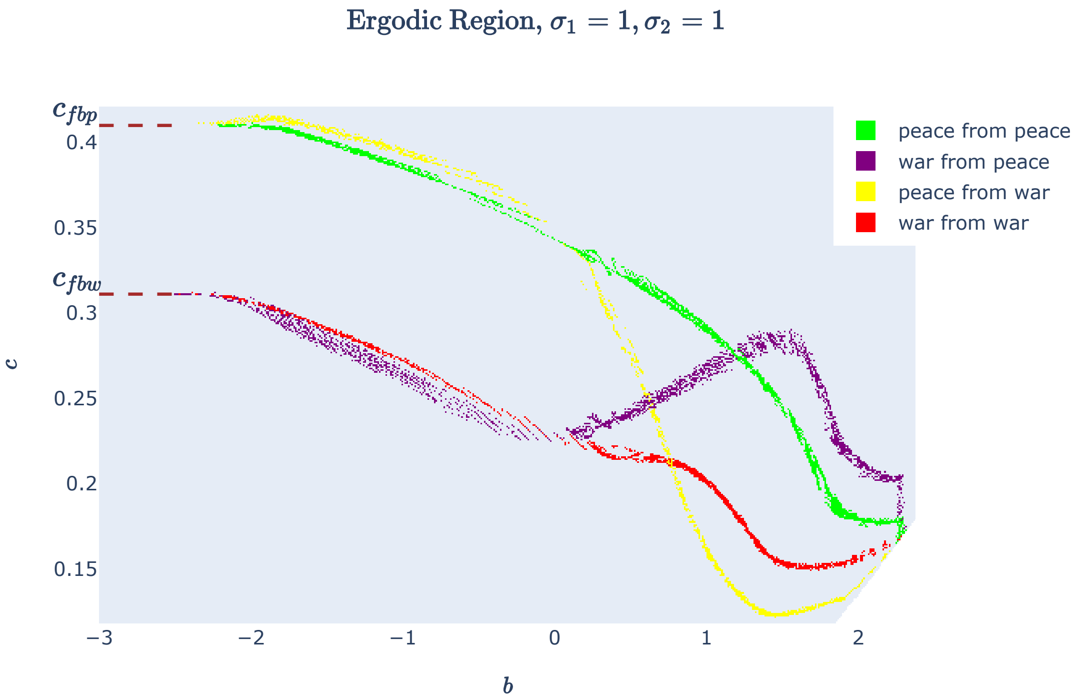

Figure 12: The four components of the ergodic set.

9.5 Policy Maps on the Ergodic Set

In dynamic economic models, we often learn much from examining the policy functions. The state is three-dimensional – b, c and gov – and all economic variables are a function of the state. However, b and c are highly correlated conditional on the war/peace state today and yesterday. In the Peace from Peace plot, for each level of b, the blue dots represent all the possible values for consumption. In general, this map from b to c is a correspondence, but so close to being a function that we can use the same visualization tools that we would use if c were a function of b. In particular, we can plot all variables as a correspondence of b and the war/peace scenario; we call these plots "policy maps".

Figure 13 displays the policy maps if the current state is peace and the previous state was also peace. Most variables depend in an intuitive manner on the current debt, b. For example, the tax rate (the scatter of green dots) tends to increase as b increases, but consumption and labor supply fall. Government spending is basically constant (close to the first best level) until b gets high at which point g has to be reduced in order to keep debt from getting too high. Note that expected consumption in the next period (the red plus signs) is essentially equal to current consumption and the interest rate is small and positive. This is consistent with the fact that the next period will almost surely be peace if we have peace today. Also note that the change in debt (the bplus-b dots) is close to zero but often slightly negative. This is all consistent with our perpetual peace plots above.

...

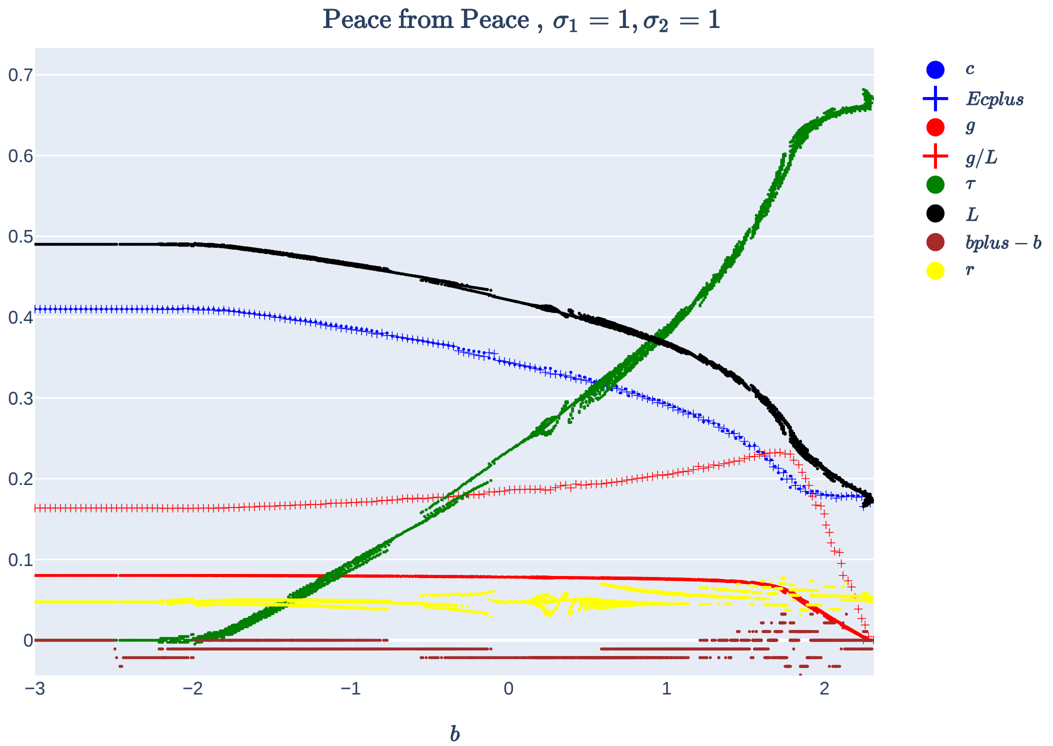

Figure 13: Policy maps in peace from peace states

Figure 14 displays the policy maps if we are in the middle of a war, which we call War from War. Again, the dependencies on b are generally intuitive. Government spending is close to the target level for negative and moderate levels of b. The tax rate rises for those levels of b because taxes have to finance both g and the cost of sustaining the debt. In fact, debt is rising quickly for these levels of b. Similarly, both consumption and labor supply are falling due to the heavy taxes and the high interest rates necessary to attract the buyers of debt. In this region, the expectation of consumption tomorrow is greater than consumption today, a natural result because there is a high (0.33) chance of the war ending.

However, for b larger than about 0.8, things change significantly. Debt accumulation continues to be positive but at a slower rate. Interest rates drop but expected future consumption is less than current consumption. a combination consistent with the purchase of new debt. At the highest levels of debt, debt accumulation ends. We see that consumption, labor supply and g are low and taxes are high. This is all consistent with our plots above of perpetual war.

The lesson from the War from War case is that economic policies are significantly affected by the ultimate level of debt even if current debt is much less than the maximum possible debt. The simple model advocated by Barro (1979) implies dynamics during a war similar to ours for negative and small b, but the dynamics for significant debt are much different.

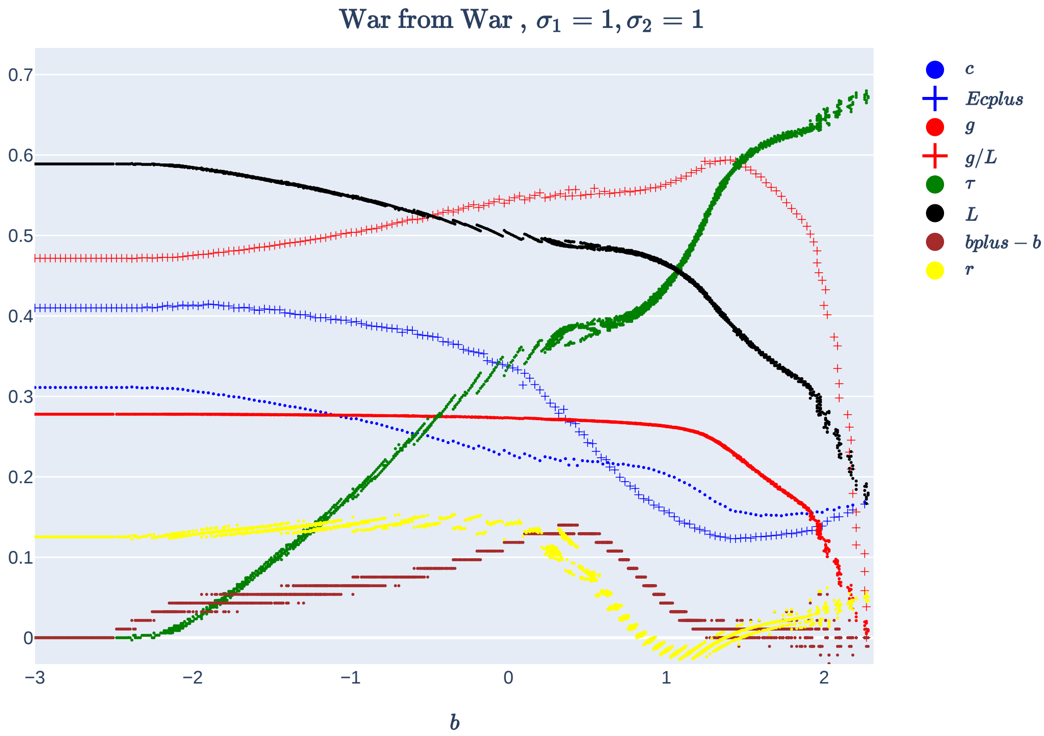

Figure 14: Policy maps in war from war states

The next plot displays the variables if the current state is war but the previous state was peace. Here we see much more complex patterns. Government spending is at the war-time first best until debt is about 0.5 after which it falls. The change in debt is positive if the current debt is less than 1. In order to sell these extra bonds, r (the yellow dots) must become high. However, if the current level of debt exceeds 1 at the beginning of a war, there will be a reduction in the level of debt, which can happen only if consumption remains high but expected consumption in the next period is less.

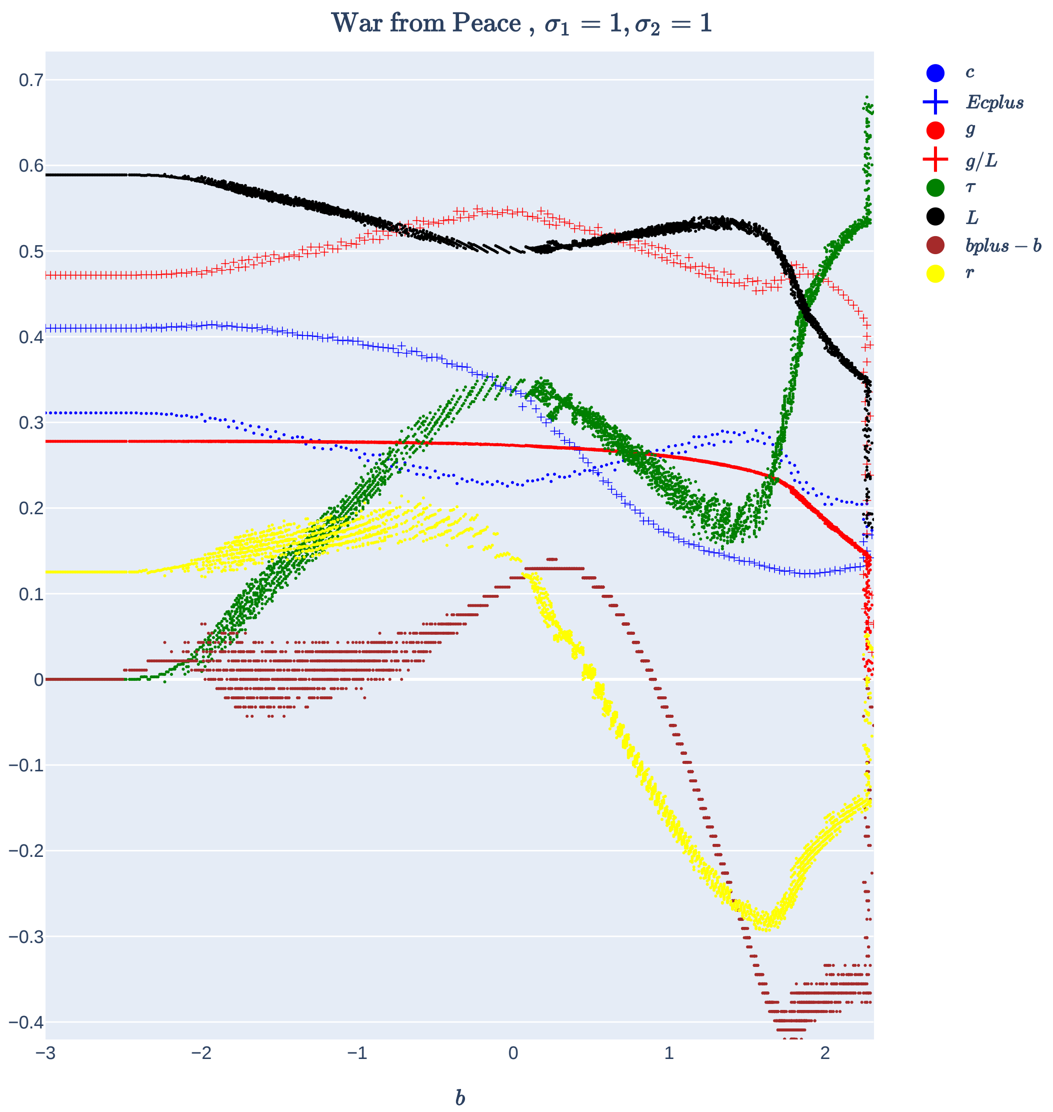

Figure 15: Policy maps in war from peace states

The last plot displays the variables if the current state is peace but the previous state was war. For small levels of debt, the policies are intuitive. Taxes increase with debt, causing consumption and labor supply to fall. g stays constant because of the strong taste to meet the target. Expected consumption growth is essentially zero. The policies are close to those in the peace-from-peace case. When debt is nontrivial, the policy responses are admittedly strange. Debt increases, which forces interest rates to rise in order for the debt to be purchased. Consumption is low, close to levels in war-to-war states, but expected growth rate is positive.

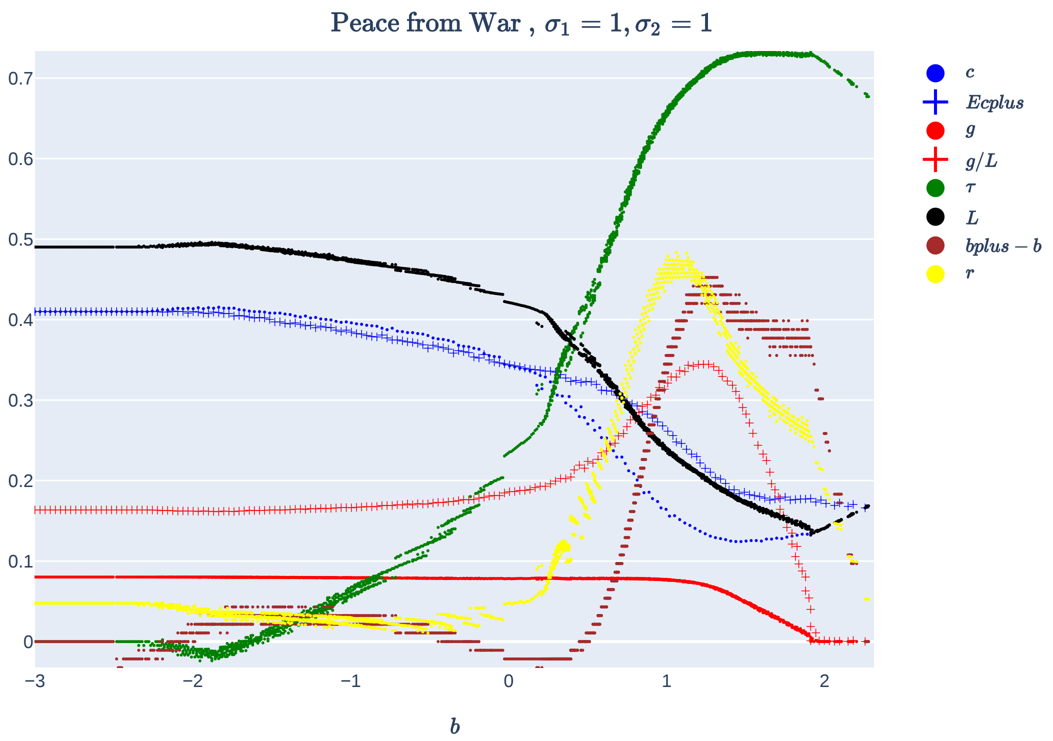

Figure 16: Policy maps in peace from war states

The main finding is that policies depend on debt in simple ways when debt is low, but the dependencies are much more complex when debt levels are high.

9.6 Simulation: distribution of debt

The next figures show how the stock of debt changes over time. The first figure displays the initial states of our simulations. At t=0, we start with 200 simulations at each chosen debt and consumption level. None of the initial states are war chest states. There are no (b,c) in the dark brown region because those states are war chest levels in both peace and war. There are some (b,c) points in the light brown region because (b,c,war) is not a war chest point. We only look at initial states in the ergodic box because the process will be in that box after one period no matter what the initial point is.

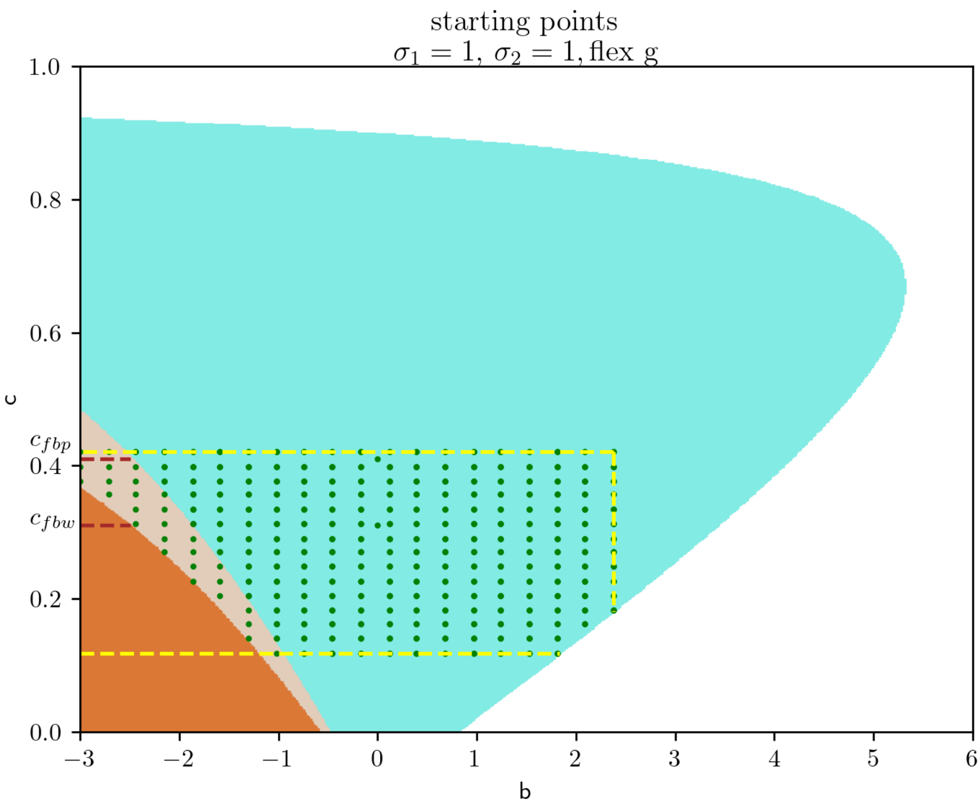

Figure 17: Add description here

The next animation displays the distribution of debt conditional on starting in one of the initial states. After one period, we see that the distribution of debt spreads out between b=-2 and b=2 for t=1, 2, 3. The ratio of blue to red increases because a third of the war states at t=0 transitioned to peace and only 2 % of the peace states transitioned to war. By t=100, we see that the distribution becomes bimodal and is increasingly bimodal a t increases. Note that the vertical scale is log. At t=50,000, most states have positive debt and looks like it is more concentrated at high debt levels. After t=1,000,000, the distribution separates into a small mass with $b < 0$ and a much larger mass around b=2. At t=2,000,000, the level of debt is almost surely at a high level. The very small number of cases with negative debt are all in the peace state. If the system is in war at t=2,000,000 then debt is around b=2. These simulations show that even if debt starts negative, war shocks will kick off a process where debt increases until it becomes large.

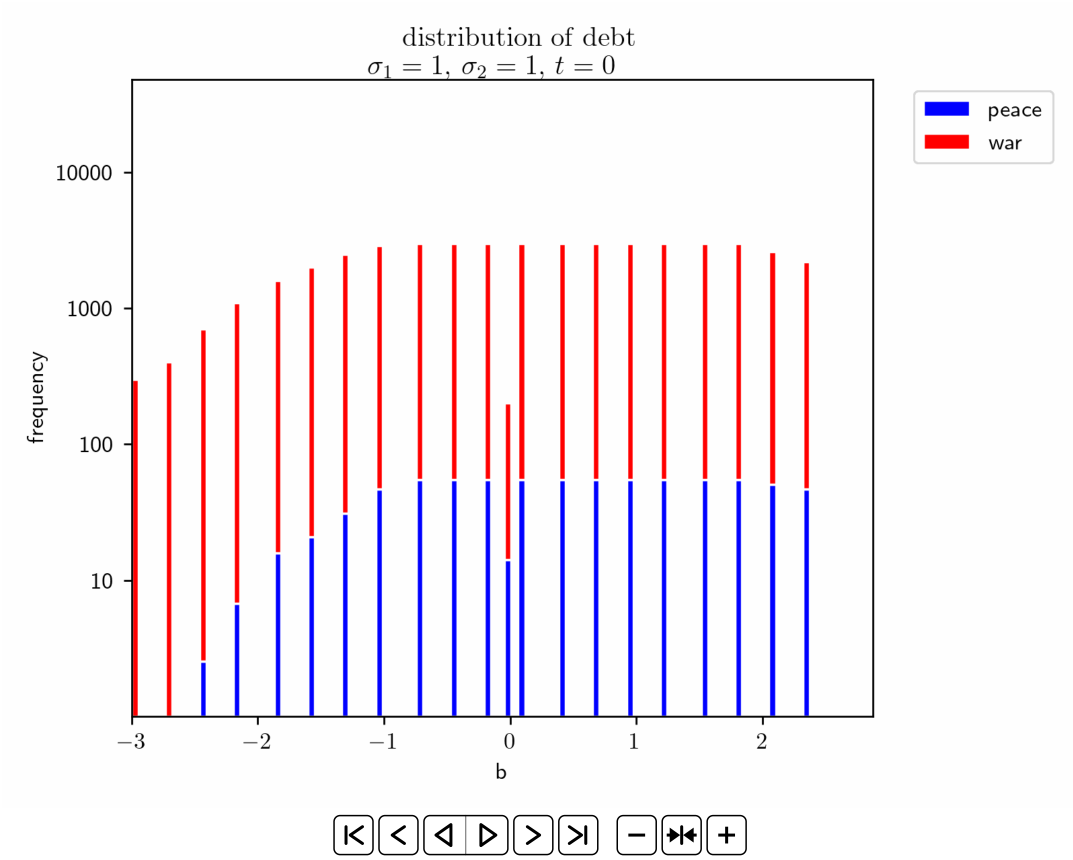

The next animation shows how consumption and debt jointly move over time. It uses the same simulations used in the previous figure. In this figure, the height is plotted on a square root scale in order to keep the display clearer. It presents a more detailed description of how a negative debt level in peacetime will evolve to a much higher debt level.. At t=0, the set of simulations begin at states that cover the ergodic set. Blue represents being initially with negative debt and red represents positive debt. At t=1, we see that the debt levels immediately move to a relatively small subset of states. We also see that most simulations with initial negative (positive) debt still have negative (positive) debt. However, by t=100, we see that some simulations that began with positive debt have moved to negative debt. After t=100, the debt level concentrates in two regions: one with low debt and high consumption and the other with high debt and low consumption. Over time, the mass of debt moves towards the high debt and low consumption region. By t=1,000,000, almost all simulations are at the high debt and low consumption points. Note that the initial level of debt does not matter, a fact represented by the one high column that has as many blue points as red.

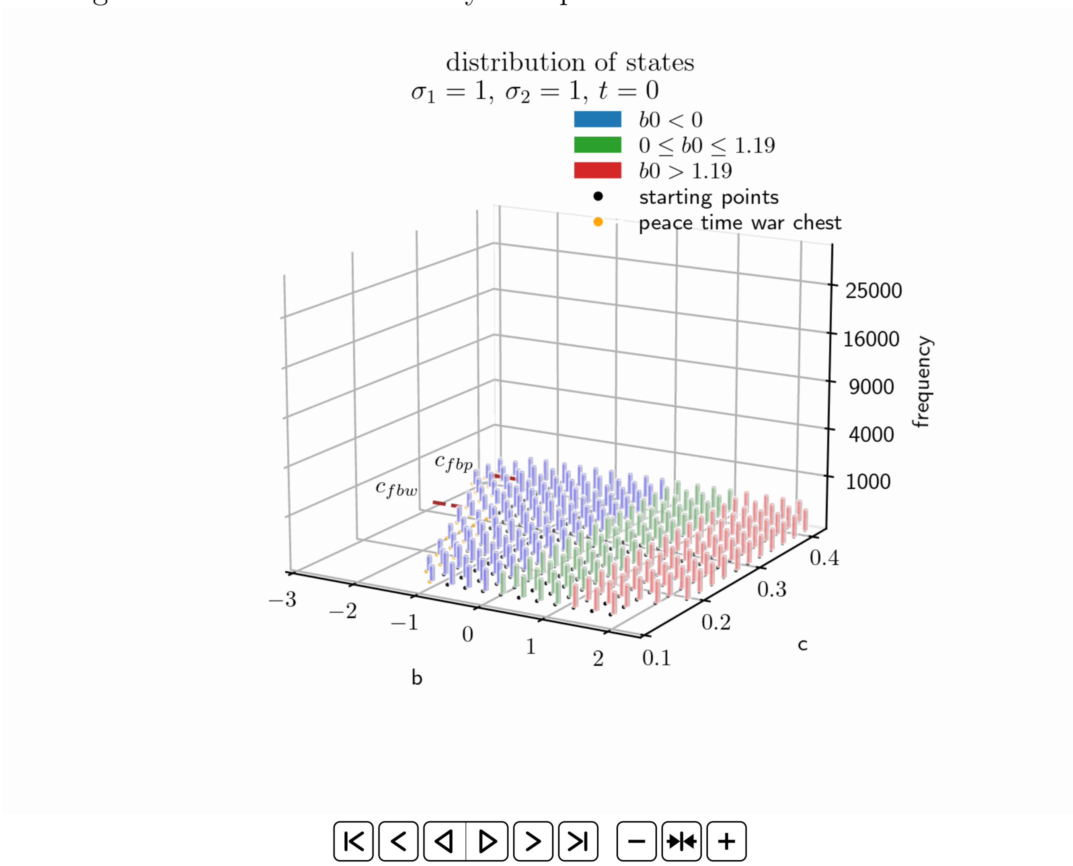

9.7 Time Series Properties of Debt and Taxes

The Barro and MSS analyses aim to justify simple time series properties for debt and taxes, arguing that debt and taxes are close to a random walk. We can estimate time series properties of debt and taxes for simulations of our model. Table 2 and table 3 displays the results from three different initial debt levels, 0.005, 1.19, and 2.14, and one initial consumption level. We distinguish between when the initial state is at peace (z0=0) and war (z0=1). coef1 is the autoregressive parameter. We see that when initial debt is low, the random walk description is valid for debt, but when initial debt is high debt has a downward drift. The results for the tax rate definitely reject the random walk hypothesis. When debt is low, the tax rate is persistent but when initial debt is high tax rate changes have low, if any, persistence.

Table 2: AR1 regression on b

σ_1	σ_2	b0	c0	z0	coef1	coef1 std	z0	coef	coef1 std
1	1	0.0051	0.4099	0	0.9914	0.0047	1	0.9710	0.0330
1	1	1.1899	0.4099	0	0.9767	0.0259	1	0.9535	0.0809
1	1	2.1377	0.4099	0	0.7728	0.1154	1	0.4429	0.2673

Table 3: AR1 regression on τ

σ_1	σ_2	b0	c0	z0	coef1	coef1 std	z0	coef1	coef1 std
1	1	0.0051	0.4099	0	0.6898	0.1029	1	0.2640	0.1240
1	1	1.1899	0.4099	0	0.5639	0.1917	1	0.1795	0.1347
1	1	2.1377	0.4099	0	0.2519	0.2067	1	-0.0142	0.0979

10 Conclusion

We can solve the model of MSS with high precision using our computational framework. When we assume exogenous spending as in MSS, we find that the feasible region is small for historically reasonable levels of government spending. In our calibration of government spending following the US experience, almost no government debt level is feasible.

In contrast, our model assumes endogenous government spending chosen according to a utility function 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 as its creditors know that spending will be cut if necessary to honor the debt.

Models like MSS and ours assume one only factor (labor), one consumption good, and identical agents are highly stylized. While 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. Therefore we do not discuss some aspects of our results, such as the initial response to a war shock, because they depend on this excessive ability of the interest to change. In future work, we plan to include capital and the ability to default.

References

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

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

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

Cai, Y., Judd, K. L., and Lontzek, T. S. (2017). The Social Cost of Carbon with Economic and Climate Risk. *The Hoover Institution Economics Working Paper Series*.

Cai, Y. and Lontzek, T. S. (2019). The Social Cost of Carbon with Economic and Climate Risks. *Journal of Political Economy*, 127(6):2684–2734.

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

Kydland, F. E. and Prescott, E. C. (1980). Dynamic optimal taxation, rational expectations and optimal control. *Journal of Economic Dynamics and Control*, 2:79–91.

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

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