Optimal Fiscal Policy with Robust Control

By

Justin Svec

October 2010

COLLEGE OF THE HOLY CROSS, DEPARTMENT OF ECONOMICS
FACULTY RESEARCH SERIES, PAPER NO. 10-04*

Department of Economics
College of the Holy Cross
Box 45A
Worcester, Massachusetts 01610
(508) 793-3362 (phone)
(508) 793-3708 (fax)

http://www.holycross.edu/departments/economics/website

*All papers in the Holy Cross Working Paper Series should be considered draft versions subject to future revision. Comments and suggestions are welcome.

Optimal Fiscal Policy with Robust Control

By

Justin Svec[†]
College of the Holy Cross

October 2010

Abstract

This paper compares the fiscal policies implemented by two types of government when confronted by consumer uncertainty. Consumers, lacking confidence in their knowledge of the stochastic environment, endogenously tilt their subjective probability model away from an approximating probability model. The government does not face this uncertainty. Through its choice of a labor tax and the supply of one-period public debt, the government manipulates the competitive equilibrium allocation and the consumers' probability distortion. I consider two types of altruistic government. A "benevolent" government maximizes the consumers' expected utility under the approximating probability model, whereas a "political" government maximizes the consumers' expected utility under the consumers' subjective probability model. I find that, relative to a full-confidence setup, the benevolent government relies more heavily on labor taxes to finance fluctuations in spending, while the political government depends more on public debt to absorb the fiscal shock. These policies are designed to re-align the consumers' savings decisions with their full-confidence values and to reduce the fluctuations in the consumers' welfare across states, respectively.

.

JEL Classification Codes: E61, E62, H21

Keywords: Robust control, uncertainty, taxes, debt, Ramsey problem

[†]Department of Economics, Box 45A, College of the Holy Cross, Worcester, MA 01610-2395, 508-793-3875 (phone), 508-793-3708 (fax), jsvec@holycross.edu

1 Introduction:

Within the optimal fiscal policy literature, consumers are typically endowed with knowledge of the probability model that characterizes the stochastic equilibrium. That is, the consumers can accurately forecast the possible state-contingent paths of the endogenous variables, which include prices and policies, as well as the exogenous variables. This ability is critical because one channel through which policy influences the equilibrium is through its effect on the consumers' expectations. By manipulating future labor taxes, for example, the government alters the consumers' expectations about the path of asset returns. These beliefs then shape the incentives faced by the consumers in earlier periods, guiding their decisions about how to allocate wealth across time and state.

The role played by fiscal policy to influence the consumers' expectations can be seen in Lucas and Stokey (1983). In this rational expectations model, it is optimal for the benevolent government to set a fairly smooth profile of labor taxes across states. This policy reduces the fluctuations in the intra-temporal distortion caused by the linear labor tax. Since taxes are less volatile across states than government spending, the government relies on public debt to finance the difference. This leads to a primary deficit when spending is high and a primary surplus when spending is low. To sustain this equilibrium, the consumers must hold the opposite profile of debt, loaning money to the government in the former case and borrowing money from the government in the latter. For the consumers to choose this pattern of savings, they must hold a particular set of beliefs about how asset returns move across state and time. Fiscal policy is designed to generate these beliefs. It does so by manipulating the stochastic discount factor, which in turn influences the consumers' forecasts of asset returns. Thus, fiscal policy's impact on consumer beliefs is an integral component of the equilibrium.

Underlying this solution is the assumption that the consumers are confident that they have the correct probability model in mind when making their decisions. However, one might worry that the equilibrium and the implied fiscal policy prescriptions hinge upon the accuracy of consumer beliefs. If, instead, consumers do not possess model-consistent expectations, they might react according to distorted forecasts of future policies and prices. This could lead to consumer behavior that undermines the government's ability to implement a smooth tax rate across states. As a result, consumer uncertainty could potentially call into question the fiscal policy prescriptions of Lucas and Stokey (1983).

The goal of this analysis, therefore, is to determine how an altruistic government responds to consumer

uncertainty. Specifically, how should a fiscal authority balance the distortions caused by the linear labor tax and consumer uncertainty? An important feature of this paper is that it analyzes two different types of altruistic government. By examining a number of different objective functions for the government, this paper attempts to disentangle the policy implications of consumer uncertainty from the planner's preferences.

As in Lucas and Stokey (1983), this paper assumes only one source of randomness: a shock to government spending. This shock will be interpreted as an extreme event, implying a large rise or fall in public expenditures[1]. The consumers and the government are both endowed with the same approximating model, which fully specifies the probabilities over all possible histories of both the endogenous and exogenous variables. The government is confident that this approximating model is an accurate description of the economy. The consumers, however, are not. The consumers are unsure about whether the approximating probability model truly characterizes the equilibrium[2].

One way to formalize the consumers' behavior given their uncertainty is through the multiplier preferences of Hansen and Sargent (2001, 2005, 2007). I will follow their formulation when developing the consumers' decision problem. The consumers, instead of trusting that the approximating model represents the truth, believe that the true probability measure lies within a range of measures. Given a finite amount of data, they worry that any probability model within this range could potentially characterize the equilibrium. The consumers respond to this type of uncertainty by applying a max-min operator to their decision problem. In doing so, the consumers endogenously distort their subjective probability model away from the approximating model[3]. This process ensures that they choose a 'robust' allocation, one that performs well even under the worst-case probability model.

It is assumed that the government is able to commit to a path of fiscal policy, chosen at time $t = 0$. Unlike the consumers, the government does not doubt its approximating model. In solving its optimization

[1] I will refer to the high government spending state as 'war' and to the low government spending state as 'peace'.

[2] There is a substantial experimental literature devoted to understanding how individuals respond to this type of uncertainty. Ellsberg (1961), for example, demonstrates that people prefer to place bets on games with known probabilities rather than unknown probabilities. These preferences suggest that people respond to uncertainty as if there was no single measure characterizing the probabilities over events. Camerer and Weber (1992) indicate that this aversion to uncertainty holds across a wide variety of environments.

[3] Importantly, this is a model of doubt, not lack of information. The consumers are aware that they are endowed with the same approximating model as the government. However, whereas the government is confident in the accuracy of this model, the consumers are not. Thus, the government has no additional information that it could reveal to the consumers in order to reduce their uncertainty.

problem, though, the government does take into account how its choice of fiscal policy affects the consumers' subjective probability model[4].

The government is altruistic and so maximizes the consumers' expected utility. When the expectations of the government and consumers coincide, there is a unique objective function for an altruistic government. When consumers face model uncertainty, though, this is no longer the case. The consumers' doubt leads them to optimize according to a subjective probability model, one that the government believes is incorrect. In this setting, there are a number of objective functions the government could have. This paper considers two types. The first type of planner considered is one that maximizes the representative consumer's expected utility under the approximating probability model. This type of government is labeled 'benevolent.' The second type of planner considered is one that maximizes the representative consumer's expected utility under the consumer's subjective probability model. This type of government is labeled 'political.'

The benevolent government represents a paternalistic planner, one that rejects the consumers' beliefs as distorted and sets policy according to what it believes the consumers should prefer. The political government, although confident in the approximating model, avoids this paternalism. Instead, the political government maximizes an objective function that is more aligned with the consumers' own preferences. Given the consumers' doubt about the correctness of the approximating probability model, the consumers might prefer this type of government. The multiplicity in planner objective functions allows me to examine the interaction between consumer uncertainty and the preferences of the government.

To foreshadow the results described below, the benevolent government relies more heavily on labor taxes to finance the shock to government spending than would be optimal if the consumers faced no uncertainty. The benevolent government chooses this volatile labor tax rate in order to reduce the distortion in the consumers' savings decisions. By increasing the labor tax during war and decreasing it during peace, the government influences asset returns, which partially re-aligns the consumers' savings decisions with their full-confidence values.

The political government, on the other hand, chooses the opposite type of policy, opting to finance more of the shock to spending through public debt. This policy smoothes the consumers' welfare across states, directly reducing the consumers' probability distortion. As these conclusions make clear, the implications of consumer uncertainty depend critically on the type of altruism of the planner.

[4]Given the path of prices and allocation, the consumers' probability distortion is fully revealed to the government.

This paper fits into a growing literature that examines whether the policy prescriptions derived from rational expectations models are robust to model uncertainty. This literature, however, largely concentrates on a different issue than the one considered in this paper. Whereas this paper analyzes the impact of consumer uncertainty, other papers in this literature generally focus on the policy implications of the government lacking confidence in the approximating probability model. For example, Dennis (2007) considers a monetary policy model in which the central bank is unsure about the stochastic process governing the shocks to the Philips' curve and the Euler equation. In addition, the central bank is also unsure about the probability model held by the firms. Given this uncertainty, a discretionary central bank reacts more aggressively to stabilize inflation than would be optimal under rational expectations.

Woodford (2010) studies a different, and novel, type of uncertainty faced by the central bank. In his model, the central bank is confident about its own model of the economy but is unsure about the beliefs entertained by the private sector. Not wanting to implement a policy that performs poorly if firms do have model-inconsistent beliefs, the bank applies a max-min operator to its decision problem. He finds that this type of uncertainty leads the central bank to restrict the degree to which cost-push shocks translate into inflation relative to a rational expectations model. Other examples in this literature include Kocherlakota and Phelan (2009) and Orphnides and Williams (2007).

One paper in the literature that discusses the impact of consumer uncertainty on policy is Karantounias, Hansen, and Sargent (2009). Independently, they also incorporate consumer uncertainty into the fiscal policy model of Lucas and Stokey (1983). The focus and scope of their analysis, however, are considerably different than mine. First, their paper only analyzes the impact of consumer uncertainty on one type of government. The goal of my analysis, though, is to compare how different types of government set policy when confronted with consumers who face uncertainty. By examining a range of preferences for the planner, this paper is better able to isolate the impact of consumer uncertainty on fiscal policy[5]. Second, Karantounias, Hansen, and Sargent (2009) focuses on whether the policy conclusions match some stylized facts of the empirical literature on US public finance, namely the persistence of fiscal policy. My analysis is more normative in approach, focusing on the incentives underlying each government's fiscal policy.

Svec (2010) further explores the impact of consumer uncertainty on optimal fiscal policy in a model

[5]In my paper, the formulation of the benevolent government's problem overlaps with that of Karantounias, Hansen, and Sargent (2007). In fact, with a redefinition of variables, my solution matches theirs. However, the focus of my analysis is to contrast the incentives of and policies chosen by different fiscal governments facing consumer uncertainty.

with capital. The government's objective function is the representative consumer's expected utility under the consumers' subjective probability measure. With these preferences, the government optimally relies more heavily on a private assets tax to finance its spending than would be optimal if consumers were confident about their probability model. In addition, the government structures the ex-post capital taxes so that the ex-ante capital tax remains quantitatively near zero. The greater volatility in the private assets tax allows the government to set a relatively smooth labor tax. This policy reduces the fluctuations in the consumers' subjective welfare, lowering their probability distortion.

The outline of the paper is as follows. Section 2 describes the structure of the economy and characterizes the representative consumer's problem. The resulting competitive equilibrium constraints hold for both types of government. Section 3 formulates the 'benevolent' government's problem and discusses the intuition behind the chosen fiscal policy. This exercise is repeated for the 'political' government in section 4. Section 5 compares the two solutions, focusing on how the incentives of each government lead to qualitatively different policy implications. Section 6 concludes.

2 The Economy:

Time is discrete in this infinite-horizon model. There are two types of agents: the government and an infinite number of identical consumers. The only source of randomness in this model is a shock to government spending, which can take on a finite number of values. Let $g^t = (g_0, ..., g_t)$ represent the history of shocks up to and including period t. The approximating probability model indicates that the probability of each history is $\pi(g^t)$. In period 0, government spending is known to be g_0 with probability 1. The government must finance its expenditure through either a linear tax on labor, τ^n, or through state-contingent, one-period debt. In each period, the government supplies the economy with a vector of these state-contingent bonds $b(g_{t+1} \mid g^t)$ at prices $p(g_{t+1} \mid g^t), \forall g_{t+1}, g^t, t \geq 0$. If $b(g_{t+1} \mid g^t)$ is held by a consumer, the government will pay out 1 unit of the consumption good if g_{t+1} occurs in the following period and zero if g_{t+1} does not occur. It is assumed that the government can commit to its history-dependent fiscal policy chosen at time 0. There is no capital in this economy.

Consumers are endowed with one unit of time each period, out of which they choose to work or enjoy leisure, $x(g^t)$. For every unit of labor supplied, one unit of output is produced. Feasible allocations must

therefore satisfy the following resource constraint:

$$c\left(g^t\right) + x\left(g^t\right) + g_t = 1 \tag{1}$$

where $c\left(g^t\right)$ denotes consumption. The consumer's wealth is composed of her after-tax labor income and the value of savings brought into that particular state. Out of her wealth, the consumer chooses an amount of consumption and savings in the state-contingent bond market.

2.1 The Consumers' Model Uncertainty:

The fundamental novelty of this model relative to Lucas and Stokey (1983) is that the consumers face model uncertainty. They are endowed with an approximating model that specifies a probability measure over future exogenous and endogenous variables. However, the consumers are uncertain whether this approximating probability model accurately characterizes the equilibrium. They fear that other probability measures could describe the stochastic nature of the economy. To ensure that these alternative probability models conform to some degree with the approximating model, restrictions must be placed on what kind of alternative models are allowed.

Following Hansen and Sargent (2005, 2007), it is assumed that each member of the set of alternative probability distributions must be absolutely continuous with respect to the approximating model. This requirement implies that the consumers only fear models that correctly put no weight on events with zero probability. That is, if fiscal policy implies that a certain event will never occur, the consumers must also believe that this is true. The type of alternative model considered, then, allows for different weights as long as the approximating model indicates that the event occurs with a weight in between zero and one. More specifically, the alternative models must be absolutely continuous over finite time intervals. This implies that the alternative models entertained by the consumers cannot be rejected with a finite amount of data, even if they could be rejected with an infinite data set. As indicated by Hansen and Sargent (2007), this restriction allows model uncertainty to have consequences for policy deep into the future.

Applying the Radon-Nikodym Theorem, there exists a measurable function, M_t, such that the subjective expectation of a random variable, X_t, can be rewritten in terms of the expectation taken with respect to the approximating probability model:

$$\widetilde{E}\left[X_t\right] = E\left[M_t X_t\right]$$

where $E[M_t] = 1$ and $\widetilde{E}$ is the subjective expectations operator. This equation allows me to reinterpret the consumers' uncertainty as uncertainty about the underlying shock process, rather than uncertainty about the distribution characterizing the policies and other endogenous variables. Consumers can be thought of as assigning the correct values to the endogenous variables for any given history of the government spending shock, even if they are uncertain about the true probability of that history occurring.

By defining an additional term, one can measure the size of the consumers' probability distortion relative to the approximating model. Let the incremental probability distortion be defined as

$$m_{t+1} = \frac{M_{t+1}}{M_t}, \forall M_t > 0$$

and $m_{t+1} = 1$ otherwise. Then, $E_t m_{t+1} = 1$. This restriction guarantees that the feared probability distributions are indeed legitimate. With this definition, the one-period distance between the alternative and approximating models is measured by relative entropy:

$$\epsilon_t(m_{t+1}) \equiv E_t m_{t+1} \log m_{t+1}$$

This measure is grounded – if $m_{t+1} = 1, \forall g_{t+1}$, then $\epsilon(m_{t+1}) = 0$ – and convex. Thus, if $\epsilon(m_{t+1})$ is small, the set of alternative models considered by the consumer is also small. As $\epsilon(m_{t+1})$ increases, the set grows larger and the consumers less confident that their approximating model governs the spending shock.

Each period's relative entropy can be aggregated and discounted to form a measure of the total distortion:

$$E_0 \sum_{t=0}^{\infty} \beta^t M_t \epsilon_t(m_{t+1})$$

This distortion measure is used in the multiplier preferences of Hansen and Sargent (2005) and characterizes how the consumers rank their allocations.

With these preferences, the consumers choose the allocation that maximizes the following objective function:

$$\min_{m_{t+1}, M_{t+1}} \sum_{t=0}^{\infty} \sum_{g^t} \beta^t \pi\left(g^t\right) M_t \left[u(c_t, x_t) + \beta\theta\epsilon_t(m_{t+1})\right]$$

The coefficient $\theta > 0$ is a penalty parameter that indicates the degree to which consumers are uncertain about the probability measure. A small θ implies that the consumers are very unsure about their approximating model, leading to large probability distortions. A high value of θ means that the consumers

have more confidence about the underlying measure, decreasing the size of the distortion. As $\theta \to \infty$, this model could collapse to the rational expectations model of Lucas and Stokey (1983).

2.2 The Consumer's Problem:

Out of her wealth, each consumer chooses how much to consume and save in state-contingent public debt. The consumer's wealth is composed of two elements: the after-tax labor income and the value of assets brought into the period. Given that the production function turns a unit of labor into one unit of output and the wage equals 1, the budget constraint in each period is

$$\sum_{g_{t+1}} p\left(g_{t+1} \mid g^t\right) b\left(g_{t+1} \mid g^t\right) + c\left(g^t\right) \leq \left(1 - \tau^n\left(g^t\right)\right)\left(1 - x\left(g^t\right)\right) + b\left(g^t\right) \tag{2}$$

Imposing the legitimacy constraint, the consumer's problem can be written recursively using the value function, $V(b, g, A)$:

$$V(b, g, A) = \max_{c,x,b'} \min_{m'} \left\{ \begin{array}{c} u(c, x) + \beta \sum_{g'} \pi\left(g' \mid g\right)\left[m' V\left(b', g', A'\right) + \theta m' \log m'\right] \\ -\lambda\left[\sum_{g'} p' b' + c - \left(1 - \tau^n\right)(1 - x) - b\right] \\ -\beta\theta\Psi\left[\sum_{g'} \pi\left(g' \mid g\right) m' - 1\right] \end{array} \right\} \tag{3}$$

where the state variable A represents the set of aggregate state variables that the consumer must track and comes from the government's problem. The consumer takes these state variables as given, believing that her decisions cannot affect their values. In tracking the aggregate state variables, the consumer is able to forecast fiscal policy after every history.

Because of the max-min operator, we must solve the consumer's problem like a Stackelberg problem, solving the inner minimization stage before we solve the outer maximization stage.

2.2.1 The Minimization Stage:

The minimization problem determines the probability distortion that minimizes the consumer's expected utility for a given allocation. There are two incentives that must be considered when finding this incremental distortion, m'. First, the incremental distortion should be distant from unity in order to lower the consumer's subjective welfare. Second, a convex penalty term penalizes the probability distortion as it diverges from one. The optimal distortion balances the marginal benefit of lowering the consumer's subjective welfare with the marginal cost due to the penalty. The optimal value of the probability distortion,

$\hat{m}'$, solves the following first order condition:

$$V\left(b', g', A'\right) + \theta\left(1 + \log \hat{m}'\right) - \theta\Psi = 0$$

Combining this condition with the additional constraint $\sum_{g'} \pi\left(g' \mid g\right) m' = 1$, we can determine the optimal probability distortion in each state in period $t+1$. Following this procedure, the optimal distortion is

$$\hat{m}' = \frac{\exp\left(\frac{-V\left(b', g', A'\right)}{\theta}\right)}{\sum_{g'} \pi\left(g' \mid g\right) \exp\left(\frac{-V\left(b', g', A'\right)}{\theta}\right)} \tag{4}$$

Equation (4) depicts the optimal tilting of the subjective probability measure away from the approximating model. The size of this tilting depends upon the consumer's subjective welfare, V, in each state in period $t+1$. If the allocation in a particular state results in a large subjective welfare relative to the average across all states, the numerator will be smaller than the denominator, meaning that $m' < 1$. As a result, the consumer places a smaller subjective weight on this state than the approximating model does. The reverse is true for an allocation that yields a small subjective welfare. Put another way, uncertainty leads each consumer to increase the subjective weight placed on low welfare states and decrease the subjective weight placed on high welfare states.

The size of the distortion also depends upon θ, the penalty parameter. A large θ decreases the probability distortion in all states in $t+1$, meaning that m' is close to 1, $\forall g'$. A small θ, conversely, implies that the probability distortions will diverge from 1, meaning that the decisions of the consumer will drastically differ from a full confidence setup.

2.2.2 The Maximization Stage:

In this stage, the consumer takes as given the prices and fiscal policy and chooses her consumption, leisure, and state-contingent bond holdings. By plugging the optimal distortion into the consumer's value function, the consumer incorporates the forecasted worst-case shock process, determined in the minimization step. The consumer then chooses an allocation, taking into account the endogeneity of the subjective probability model. The resulting recursive problem is

$$V\left(b, g, A\right) = \max_{c, x, b'} \left\{ \begin{array}{c} u\left(c, x\right) - \beta\theta \log \sum_{g'} \pi\left(g' \mid g\right) \exp\left(\frac{-V\left(b', g', A'\right)}{\theta}\right) \\ -\lambda\left[\sum_{g'} p' b' + c - \left(1 - \tau^n\right)\left(1 - x\right) - b\right] \end{array} \right\} \tag{5}$$

The consumer's first order conditions with respect to c, x, and b' are

$$c : u_c(c, x) - \lambda = 0 \tag{6}$$

$$x : u_x(c, x) - \lambda(1 - \tau^n) = 0 \tag{7}$$

$$b' : \frac{\beta \pi(g' \mid g) V_b(b', g', A') \exp\left(\frac{-V(b', g', A')}{\theta}\right)}{\sum_{g'} \pi(g' \mid g) \exp\left(\frac{-V(b', g', A')}{\theta}\right)} - \lambda p' = 0 \tag{8}$$

and the envelope condition is

$$V_b(b, g, A) = \lambda$$

As is standard in models that assume the government has access to a distortionary labor tax, the first order conditions imply the following intra-temporal tradeoff between consumption and leisure:

$$\frac{u_x(c, x)}{u_c(c, x)} = 1 - \tau^n \tag{9}$$

A larger tax increases the intra-temporal wedge. The Euler equation when consumers face model uncertainty is

$$u_c(c, x) p' = \beta \pi(g' \mid g) u_c(c', x') m' \tag{10}$$

The consumer's fears influence her expected future marginal utility of consumption. For a given price, the consumer will choose different path of consumption, savings, and leisure than if she were fully confident in the approximating probability model.

Given these conditions, I can now define a competitive equilibrium:

Definition 1 *A competitive equilibrium is an allocation* $\left\{c\left(g^t\right), x\left(g^t\right), V\left(g^t\right), b\left(g^{t+1}\right)\right\}_{t=0}^{\infty}$, *probability distortions* $\left\{m\left(g^{t+1}\right), M\left(g^{t+1}\right)\right\}_{t=0}^{\infty}$, *prices* $\left\{p\left(g^{t+1}\right)\right\}_{t=0}^{\infty}$, *and policies* $\left\{\tau^n\left(g^t\right)\right\}_{t=0}^{\infty}$ *such that*

1. *Given the consumer's allocation, the probability distortion* $\left\{m\left(g^{t+1}\right), M\left(g^{t+1}\right)\right\}_{t=0}^{\infty}$ *solves the consumer's minimization problem,*

2. *Given the government's policy and prices, the allocation* $\left\{c\left(g^t\right), x\left(g^t\right), b\left(g^{t+1}\right)\right\}_{t=0}^{\infty}$ *solves the consumer's maximization problem, forecasting the response of the malevolent agent, and*

3. *All markets clear.*

3 The Planner's Problem

With the competitive equilibrium defined, I can now discuss the planner's problem. The planner's problem is written in its primal representation. This formulation allows the government to directly choose the representative consumer's allocation, taking into account how the consumers' subjective probability model evolves. Given this choice, the competitive equilibrium constraints then determine the necessary prices and policies that support the allocation and distortions. It is assumed that the government is able to commit to this fiscal policy at time 0.

In the following sections, I consider two types of government altruism. Given each objective function, the government chooses the competitive equilibrium that maximizes its preferences. With the resulting allocation and distortions, I back out the prices and policies that support the competitive equilibrium. A discussion of the results then follows.

3.1 The Benevolent Government:

It is assumed that the objective function for the benevolent government is the consumers' expected utility under the approximating probability model.

Definition 2 *The Ramsey problem of the benevolent government is to choose the competitive equilibrium that maximizes the expected utility of the representative consumer under the approximating model. The Ramsey outcome under the benevolent government is the competitive equilibrium that attains the maximum.*

Proposition 1 *The allocation and distortions in the Ramsey outcome under the benevolent government solve the following problem:*

$$\max_{c_t, x_t, V_t, b_{t+1}, m_{t+1}} \sum_{t=0}^{\infty} \sum_{g^t} \beta^t \pi\left(g^t\right) u\left(c_t, x_t\right)$$

subject to

$$\beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) u_c\left(c_{t+1}, x_{t+1}\right) m_{t+1} b_{t+1} + u_c\left(c_t, x_t\right)\left(c_t - b_t\right) - u_x\left(c_t, x_t\right)\left(1 - x_t\right) = 0 \tag{11}$$

$$m_{t+1} = \frac{\exp\left(\frac{-V_{t+1}}{\theta}\right)}{\sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) \exp\left(\frac{-V_{t+1}}{\theta}\right)} \tag{12}$$

$$V_t = u\left(c_t, x_t\right) + \beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) \left\{m_{t+1} V_{t+1} + \theta m_{t+1} \ln m_{t+1}\right\} \tag{13}$$

$$c_t + x_t + g_t = 1 \tag{14}$$

Proof. When setting its policy, the government is restricted in the set of feasible allocations that it can achieve by the competitive equilibrium constraints. The claim is that those restrictions are summarized by the constraints $(11)-(14)$. To demonstrate this, I will first show that any allocation and probability distortion that satisfies the competitive equilibrium constraints must also satisfy $(11)-(14)$. (2) holds with equality in equilibrium. Insert (6), (7), and (8) into (2) to get (11). (12) follows directly from the optimality condition in the inner minimization, (13) is the consumer's Bellman equation, and (14) is the resource constraint. Thus, $(11)-(14)$ are necessary conditions that the Ramsey outcome must solve. Going in the other direction, given an allocation and distortions that satisfy $(11)-(14)$, policies and prices can be determined from the representative consumer's first order conditions. ■

The first constraint, a period implementability constraint, depicts the transition equation of public debt. This equation is similar to the constraint that arises when consumers do not face model uncertainty, except that the expectation of tomorrow's value of debt is tilted by the consumers' probability distortion.

The second implementability constraint (13) captures how the consumers' value function evolves across time and states. The planner must keep track of the consumers' value function in order to take into account their probability tilting. This equation is a new constraint that does not appear in the full confidence framework. The planner also faces the resource constraint and the description of the optimal probability distortion. The constraints $\{(11),(12),(13),\text{ and }(14)\}$ fully characterize the set of competitive equilibrium restrictions.

3.1.1 Sequential Formulation of the Benevolent Planner's Problem:

The benevolent government's optimization problem is

$$\mathcal{L}=\sum_{t=0}^{\infty}\beta^{t}\sum_{g^{t}}\pi\left(g^{t}\right)\left\{\begin{array}{c} u\left(c_{t},x_{t}\right)+\mu_{t}\left[c_{t}+x_{t}+g_{t}-1\right] \\ +\xi_{t}\left[\beta\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^{t}\right)u_{c}\left(c_{t+1},x_{t+1}\right)m_{t+1}b_{t+1}+u_{c}\left(c_{t},x_{t}\right)\left(c_{t}-b_{t}\right)-u_{x}\left(c_{t},x_{t}\right)\left(1-x_{t}\right)\right] \\ +\Gamma_{t}\left[V_{t}-u\left(c_{t},x_{t}\right)-\beta\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^{t}\right)\left\{m_{t+1}V_{t+1}+\theta m_{t+1}\ln m_{t+1}\right\}\right] \\ +\beta\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^{t}\right)\varpi_{t+1}\left[m_{t+1}-\frac{\exp\left(\frac{-V_{t+1}}{\theta}\right)}{\sum_{g_{t+1}}\pi(g_{t+1}|g^{t})\exp\left(\frac{-V_{t+1}}{\theta}\right)}\right] \end{array}\right\}$$

The first order conditions are

$$\begin{aligned} c_t, \forall t \quad &\geq \quad 1: \\ 0 \quad &= \quad u_c(c_t, x_t) + \mu_t - \Gamma_t u_c(c_t, x_t) + \xi_{t-1} u_{cc}(c_t, x_t) m_t b_t \\ &\quad + \xi_t \left[u_{cc}(c_t, x_t)(c_t - b_t) + u_c(c_t, x_t) - u_{cx}(c_t, x_t)(1 - x_t) \right] \end{aligned} \tag{15}$$

$$\begin{aligned} x_t, \forall t \quad &\geq \quad 1: \\ 0 \quad &= \quad u_x(c_t, x_t) + \mu_t - \Gamma_t u_x(c_t, x_t) + \xi_{t-1} u_{cx}(c_t, x_t) m_t b_t \\ &\quad + \xi_t \left[u_{cx}(c_t, x_t)(c_t - b_t) + u_x(c_t, x_t) - u_{xx}(c_t, x_t)(1 - x_t) \right] \end{aligned} \tag{16}$$

$$\begin{aligned} V_t, \forall t \quad &\geq \quad 1: \\ 0 \quad &= \quad \Gamma_t - \Gamma_{t-1} m_t + \left(\frac{1}{\theta} \right) m_t \left[\varpi_t - E_{t-1} m_t \varpi_t \right] \end{aligned} \tag{17}$$

$$\begin{aligned} m_{t+1}, \forall g^{t+1}, \forall t \quad &\geq \quad 0: \\ 0 \quad &= \quad \xi_t u_c(c_{t+1}, x_{t+1}) b_{t+1} - \Gamma_t \left[V_{t+1} + \theta (1 + \ln m_{t+1}) \right] + \varpi_{t+1} \end{aligned} \tag{18}$$

$$\begin{aligned} b_{t+1}, \forall g^{t+1}, \forall t \quad &\geq \quad 0: \\ 0 \quad &= \quad \xi_t m_{t+1} - \xi_{t+1} \end{aligned} \tag{19}$$

The first order conditions (17) and (19) imply that both Lagrange multipliers on the implementability constraints (ξ_t, Γ_t) are martingales under the approximating model:

$$\begin{aligned} E_t \xi_{t+1} &= \xi_t \\ E_{t-1} \Gamma_t &= \Gamma_{t-1} \end{aligned}$$

This result implies that the allocation exhibits persistence. This persistence appears because the benevolent government must take into account the endogeneity of the consumers' probability distortion, which is itself a martingale. Only as $\theta \to \infty$ do the Lagrange multipliers reduce to constants, eliminating the persistence. As a result of this persistence, consumer uncertainty leads to variations in the shadow values of the marginal utility of debt and the consumers' welfare across time and states. This variation is absent in Lucas and Stokey (1983).

Beyond the inter-temporal persistence, model uncertainty also imparts additional intra-temporal smoothing. Whereas the implementability constraint is the sole condition that links the allocation across states in Lucas and Stokey (1983), the probability distortion directly connects the allocation across states. The linkage is most salient in (17), as the movement of Γ_t depends upon the difference between one state's characteristics and its expectation. Again, only as $\theta \to \infty$ does the additional intra-temporal connection disappear.

The time 0 first order conditions are

$$\begin{aligned} c_0 \quad : \quad & 0 = u_c(c_0, x_0) + \mu_0 - \Gamma_0 u_c(c_0, x_0) \\ & + \xi_0 \left[u_{cc}(c_0, x_0)(c_0 - b_0) + u_c(c_0, x_0) - u_{cx}(c_0, x_0)(1 - x_0) \right] \end{aligned}$$

$$\begin{aligned} x_0 \quad : \quad & 0 = u_x(c_0, x_0) + \mu_0 - \Gamma_0 u_x(c_0, x_0) \\ & + \xi_0 \left[u_{cx}(c_0, x_0)(c_0 - b_0) + u_x(c_0, x_0) - u_{xx}(c_0, x_0)(1 - x_0) \right] \end{aligned}$$

$$V_0(g_0) : \Gamma_0 = 0$$

In order to determine the specific values of the allocation, prices, policies, and probability distortions, I must numerically solve this model. Consequently, I formulate the recursive problem of the government below. This recursive problem uses the fact that the solution is recursive in the Lagrange multipliers on the implementability constraints to determine which variables should be added to the list of state variables.

3.1.2 Recursive Formulation of the Benevolent Planner's Problem:

In deriving the recursive form of the government's optimization problem, I assume that government expenditures follow a Markov process with transition matrix Π. Due to the time-inconsistency of the planner's problem, I apply the Marcet and Marimon (1998) procedure to the implementability constraints. The co-state variable on (11) is ξ_- with a state-contingent increment of ξ_g. The subscript g implies a state-contingent value in period $t \geq 1$. The co-state variable on (13) is Γ_- with a state-contingent increment of Γ_g. An ex-ante value function is necessary to account for probability distortion, which connects all states in a particular period.

The planner's problem in recursive form is

$$W\left(\xi_{-},\Gamma_{-},g_{-}\right)=\min_{\xi_g,\Gamma_g}\max_{c_g,x_g,m_g,V_g,b_g}\sum_{g}\pi\left(g\mid g_{-}\right)\left\{\begin{array}{c}u\left(c_g,x_g\right)+\mu_g\left[c_g+x_g+g-1\right]+\xi_{-}\left[u_c\left(c_g,x_g\right)m_gb_g\right]\\+\xi_g\left[u_c\left(c_g,x_g\right)\left(c_g-b_g\right)-u_x\left(c_g,x_g\right)\left(1-x_g\right)\right]\\-\Gamma_{-}\left[m_gV_g+\theta m_g\ln m_g\right]+\Gamma_g\left[V_g-u\left(c_g,x_g\right)\right]\\+\varpi_g\left[m_g-\frac{\exp\left(\frac{-V_g}{\theta}\right)}{\sum_{g}\pi(g|g_{-})\exp\left(\frac{-V_g}{\theta}\right)}\right]+\beta W\left(\xi_g,\Gamma_g,g\right)\end{array}\right\}$$

The initial values of the co-state variables are 0, representing the assumption that the planner at time t=0 is not bound by any previous promises. The first order and envelope conditions from this problem are described in Appendix A.

3.1.3 Model Solution and Discussion:

In order to understand how consumer uncertainty affects optimal fiscal policy, this section compares the rational expectations equilibrium with the equilibrium under consumer uncertainty. To ease the exposition, assume a simple process for government spending:

$$\begin{aligned}g_t &= 0,\forall t\neq T\\ g_T &= \left\{\begin{array}{c}G,\text{with probability }\alpha\\0,\text{with probability }1-\alpha\end{array}\right\}\end{aligned}$$

Additionally, $b_0=0$, so that consumers have no debt or assets in the initial period.

I will first discuss the rational expectations solution. As shown in Lucas and Stokey (1983), the allocation depends only upon the current value of the government spending shock. This means that there are two possible values for consumption and leisure: $\{c\left(0\right),x\left(0\right)\}$ and $\{c\left(G\right),x\left(G\right)\}$. Using (9), the tax rate also takes on two values: $\tau^n\left(0\right)$ and $\tau^n\left(G\right)$. As government spending is zero for all periods except T, the government's positive tax $\tau^n\left(0\right)$ yields a surplus at each of these dates. The surplus, equal to $\tau^n\left(0\right)\left(1-x\left(0\right)\right)$, is loaned to the consumers in each period $t<T$ at an interest rate of $\frac{1}{\beta}$, in addition to the accumulated assets from the previous period.

In period T-1, the government uses its accumulated assets to obtain insurance from the consumers against the shock at T. The government buys bonds (setting $b_T\left(G\right)<0$) at a price $\beta\alpha\frac{u_{c(c(G),x(G))}}{u_{c(c(0),x(0))}}$ that can be redeemed if $g_T=G$. In addition to using its accumulated wealth, the government finances the cost of this insurance through period T-1's primary surplus and through issuing debt that pays off if $g_T=0$. The price of this debt is $\beta\left(1-\alpha\right)$.

In period T, the allocation depends upon whether or not the shock occurs. If $g_T = G$, the government runs a primary deficit $\tau^n(G)(1 - x(G)) - G < 0$. This deficit is financed partially by the interest repayment on the loan made in the previous period and partially by selling additional debt to the consumers. For each period $t > T$, the government collects tax revenues and pays off previous debt through its primary surplus. If $g_T = 0$, though, the government uses its primary surplus and additional borrowed money from the consumers to finance its previous debt from that period forward.

Viewing this result from an insurance perspective, the government's chosen profile of tax and debt mitigates the cost of the spending shock by spreading the intra-temporal distortion across time and states. Rather than having no distortion when $g_t = 0$ and a large distortion when $g_T = G$, the government charges a positive tax on labor income in all periods leading up to T. In period T-1, these assets are then used to buy insurance from the consumers, who agree to pay a fraction of the cost of the spending shock if it occurs. The additional funds come from tax revenue and from borrowing money from the consumers in period T. Thus, under rational expectations, a primary role of state-contingent debt is to provide insurance, allowing the government to reduce its dependence on the linear labor tax to finance the spending shock.

Model uncertainty complicates this simple relationship. In addition to the insurance incentive depicted above, the benevolent government also seeks to mitigate the impact of the consumers' probability distortion. The government therefore must use fiscal policy to balance the social costs stemming from both the linear labor tax and from model uncertainty.

To see how these tradeoffs are balanced, I have computed the solution to the benevolent government's problem and have graphed the solutions below. In calculating these solutions, I have assumed that the consumers have the following CRRA preferences:

$$u(c, x) = \frac{c^{1-\sigma}}{1-\sigma} + \eta \frac{x^{1-\delta}}{1-\delta}.$$

Government spending follows the process:

$$g_t = \bar{g} + \rho\left(g_{t-1} - \bar{g}\right) + \epsilon.$$

Depending on the value of ρ, this process could resemble an iid shock to government spending or an AR(1) process. Shocks to government spending will be drawn from an approximation to a normal distribution, $\epsilon \sim N\left(0, 0.02^2\right)$. In the numerical calculations, three values of the shock are considered. The probability

of being hit by the high spending shock is 17%, which, because of symmetry, is also the probability of being hit by the low shock. In the graphs below, I plot only the solutions associated with a high value of spending (labeled "war") and a low value of spending (labeled "peace"). Consumers begin with no assets: $b_0 = 0$. The parameters used in these calculations are listed below in Table 1:

Parameters:	
Utility:	Government Spending:
$\sigma = 2$	$\bar{g} = 0.1$
$\eta = 1$	$\rho = 0$
$\delta = 2$	$g_0 = \bar{g}$
$\beta = 0.95$	

Table 1: Parameter values

An important feature of the numerical calculations is to define the true probability model. The government is confident that the approximating probability model is correct, while the consumers worry about a range of alternative probability models. Depending on what model is chosen, the government's confidence could be well-placed or the consumers' fear could be justified. For numerical simplicity, this paper assumes that the approximating probability model happens to be correct. This fact will tilt the welfare results in favor of the benevolent government, as its objective function uses the approximating probability model.

To understand how the solution changes relative to the Lucas and Stokey (1983) benchmark, I have plotted the solutions as a function of the level of consumer uncertainty. When log(theta) is large, the consumers are confident in the approximating model, and the equilibrium approaches the rational expectations solution discussed above. As log(theta) falls, however, the consumers are increasingly uncertain about the probability model. Because this larger uncertainty leads to larger behavioral distortions (from the vantage point of the benevolent government), the government becomes more concerned about the costs stemming from model uncertainty.

As analytically shown above, the consumers' uncertainty leads them to tilt their subjective probabilities away from the approximating model. This probability distortion can be seen in Figure 1. As log(theta) falls, the consumers increase the weight placed on the high government spending state and decrease the weight placed on the low government spending state. Thus, the consumers worry that war is more likely

and peace is less likely than indicated by the approximating probability model.

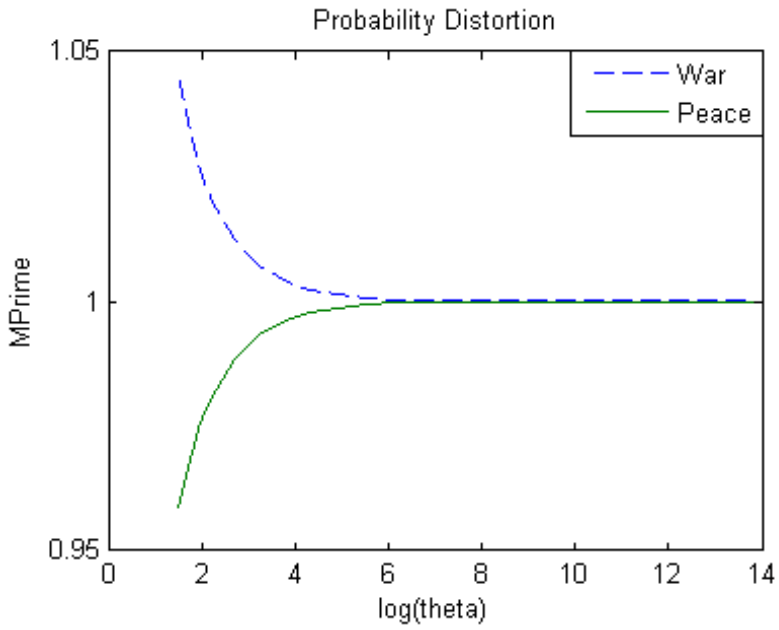

Figure 1: Consumers' incremental probability distortion for different levels of θ

One important consequence of this probability tilting is that the consumers choose a different profile of savings than would be optimal if they were confident in the approximating probability model. For the same price level, because the consumers place a higher subjective probability on war, they would like to increase their holdings of the war-contingent asset. Similarly, because the consumers place less weight on the state of peace, they desire fewer of these state-contingent assets. These shifts in demand, in return, have implications for the prices and returns of the state-contingent assets. As seen in the pricing equation, (10), the increase in demand for the war-contingent bond raises the price of the war-contingent bond, while the decrease in demand for the peace-contingent bond lowers the price of the peace-contingent bond. As a result, the returns on the war-contingent bond falls and the peace-contingent bond rises.

Critically, this savings profile undermines the government's ability to obtain insurance against its spending shock. In the full-confidence framework, the benevolent government would like to buy war-contingent debt (to be paid off by the consumers in the event of war) and sell peace contingent debt (to be paid off by the government in the event of peace). This profile enables the government to set a smooth labor tax rate. However, as indicated above, model uncertainty leads consumers to desire the exact opposite profile of debt. The consumers, in their uncertainty about the true probability distribution, want to save in a manner that makes it difficult for the government to use public debt as insurance against

the fiscal spending shock.

Given this tension, how does a benevolent government resolve the competing uses for public debt? The solution hinges on the fact that the benevolent government's objective is to maximize the consumers' expected utility under the approximating probability model. Because the benevolent government optimizes with respect to this model, it distrusts the consumers' subjective probability model and so believes that the consumers are distorting their savings profile. Consequently, the benevolent government attempts to use its fiscal policy instruments to re-align the consumers' savings decisions with those that would be optimal under the approximating probability model. It accomplishes this by manipulating the prices and returns on debt, as seen in Figure 2. Specifically, fiscal policy is designed to raise (lower) the price of war-contingent (peace-contingent) assets. The higher price on war-contingent assets discourages the consumers from holding this debt instrument. This higher price also allows the government to increase the amount of money it loans to the consumers that they must repay during times of war. This represents the insurance that the government gains as a result of its policy. Conversely, the lower price on peace-contingent assets encourages the consumers to loan money to the government that it must repay during times of peace. This represents the premium that the government must pay for the insurance.

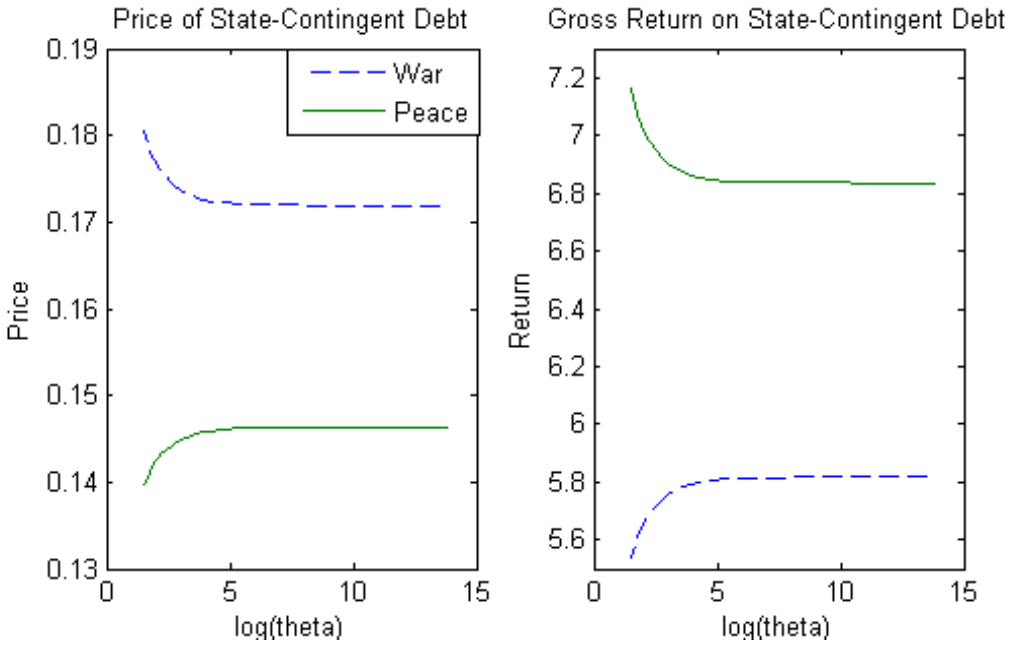

Figure 2: Price and return on debt for different levels of θ

To be clear, two factors drive the movement in state-contingent asset prices. The first factor is that the consumers' subjective probability model increases the demand for war-contingent bonds and decreases

the demand for peace-contingent bonds. This puts upward pressure on the price of war-contingent assets and downward pressure on the price of peace contingent assets. The second factor is that the government implements a policy that intensifies these price movements, driving the prices on war-contingent (peace-contingent) assets even higher (lower). These two factors are decomposed in Figure 3. This graph plots the war-contingent asset price under two different scenarios. The solid line, labeled "Beliefs", depicts the movement in the price due purely to the distorted beliefs of the consumers. This line is obtained by multiplying the stochastic discount factor that arises when consumers are confident in the probability model with their probability distortion across different levels of θ[6]. The dotted line, labeled "Policy", depicts the price of the war-contingent asset, taking into account both the consumers' beliefs and the fiscal policy.

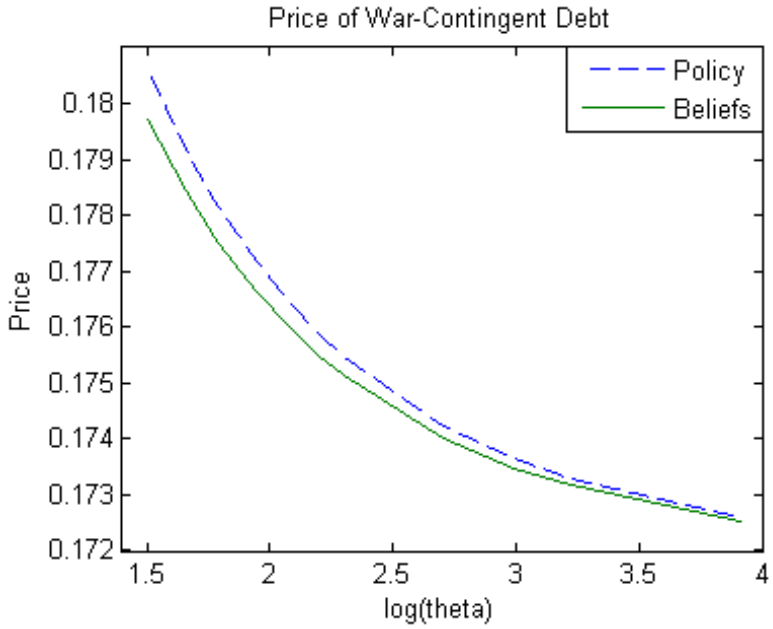

Figure 3: Decomposing price movements into beliefs and policy

The question remains as to how the government uses the state-contingent labor tax rate to accomplish the price movements discussed above. Relative to the policy chosen when consumers face no uncertainty, model uncertainty leads the government to implement a relatively high (low) labor tax rate conditional on war (peace), as seen in Figure 4. This policy affects the allocation chosen by the consumers in each

[6]This relative contribution of beliefs in price movements is only approximate. By keeping the stochastic discount factor constant at the $\theta \to \infty$ value, I am assuming that the allocation remains at its "certainty" level even though the consumers' uncertainty is growing. However, the probability distortions and the allocation move together, making it difficult to separate the role of beliefs and policy. This exercise is merely meant to hint at the two factors that drive price movements.

state. As a result, the stochastic discount factor moves in such a way as to raise the war-contingent bond price and lower the peace-contingent bond price.

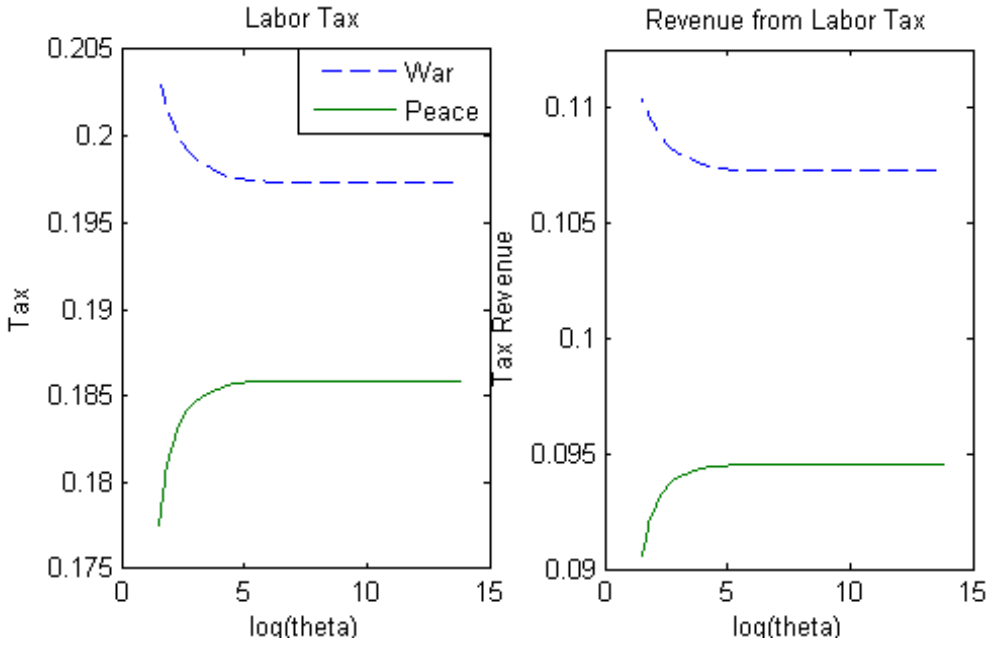

Figure 4: Optimal fiscal policy implemented by the benevolent government for different levels of θ

To elucidate this point, consider the increase in the war-time tax in T+1. The higher tax encourages consumers to enjoy more leisure during wars because the after-tax marginal return on labor has fallen. Due to the decrease in labor income, consumers reduce their consumption. This change then affects the price of the war-contingent debt in period T. For a given marginal utility of consumption at T, the increase in the labor tax rate in war will raise the marginal utility of consumption at T+1, causing the price of war-contingent debt to rise. Consumers, faced with the increased price and lower return on war-contingent debt, choose to hold less of this debt than if the labor tax had not increased. The opposite profile of incentives holds true in the event of peace. In effect, the government sets fiscal policy to discourage consumers from holding war-contingent debt and from borrowing peace-contingent debt, partially reversing the savings distortion.

The changes in the labor tax rate have important macroeconomic implications. First, the fluctuations in the primary deficit are less pronounced than what would be optimal if consumers were confident in their knowledge of the stochastic environment. Second, relative to the full-confidence framework, the consumers decrease their labor supply during times of war and increase it during times of peace. This has the implication of reducing the volatility of output across state. Third, the movement in the consumers' consumption matches that of their labor supply: consumption is lower during war and higher during peace

relative to the case in which consumers face no model uncertainty.

The conclusions described above characterize the volatility of the solutions at a particular point in time. Below, I compare the impulse response functions under different degrees of model uncertainty. To create these response functions, I assume that government spending is equal to its average, $g_t = \bar{g}$, for all periods except period 1, at which time $g_1 = g_{high}$. The spending profile associated with this one-period war can be seen in the top left graph in Figure 5.

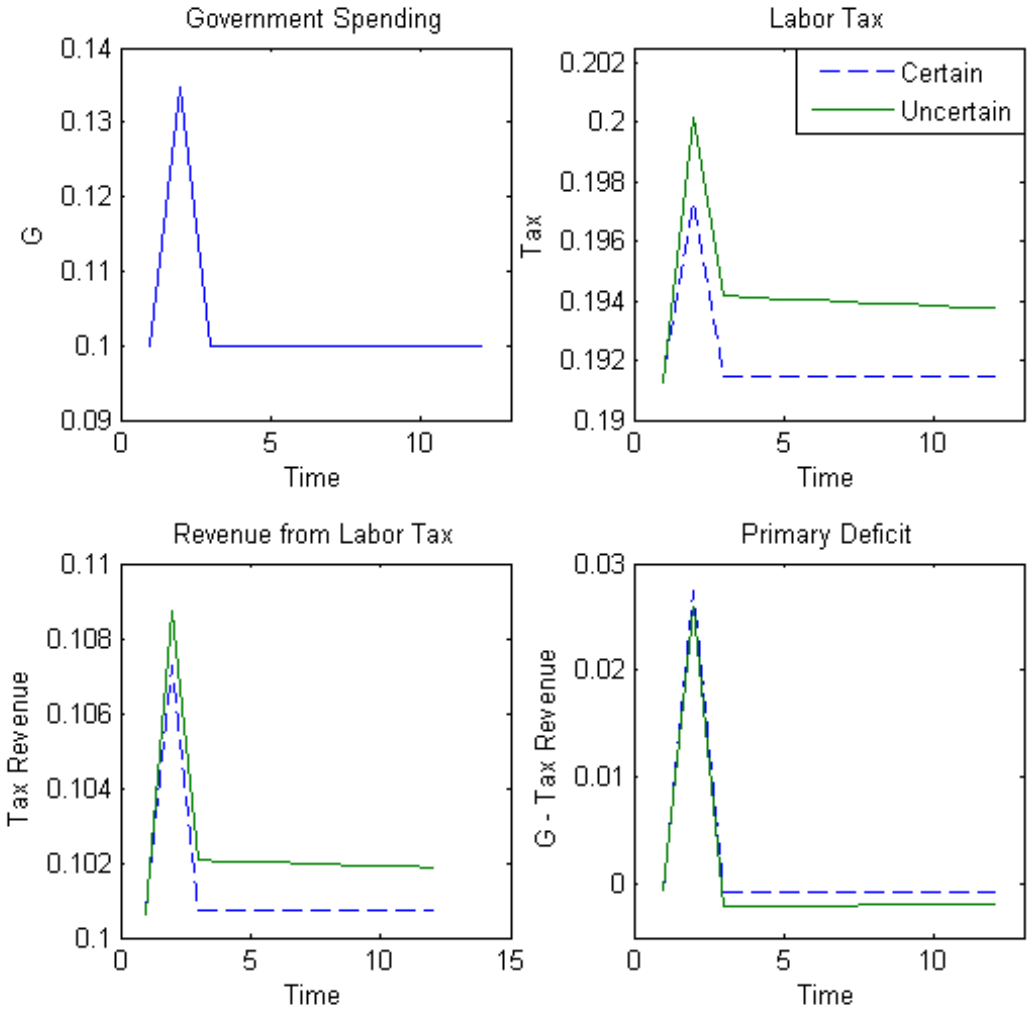

Figure 5: Comparing the policy impulse response functions to an increase in government spending when $\theta \to \infty$ and $\theta = 10$

During periods of war, the benevolent government raises the labor tax rate to help finance the spending shock. Relative to the case in which consumers completely understand the shock process, model uncer-

tainty leads to a higher spike in the tax rate. This increase, although partially offset by a decrease in labor supply, results in a rise in labor tax revenues. As indicated above, the benevolent government relies more heavily on labor taxes to finance the spending shock when consumers face model uncertainty.

An additional feature of these impulse response functions that is worthy of note is that the levels of persistence differ across the two models. When consumers are certain, the post-war policy returns to its pre-war levels after the shock ends. This is not the case when consumers face uncertainty. Instead, the war-time labor tax remains higher than its full confidence value for a number of periods after the one-period war. This persistence translates into a prolonged period of decreased consumption and labor hours for the consumers.

3.2 The Political Government:

An implication of model uncertainty is that the consumers' subjective expectation could differ from the true expectation. This distinction leads to some flexibility as to the objective function of an altruistic government. The previous section modeled a benevolent government that maximizes the consumers' expected utility under the approximating probability model. This objective function leads the government to choose a volatile tax rate that is meant to manipulate the consumers' expectations about bond returns. This section describes the decision problem of a political government, which maximizes the consumers' expected utility under their own subjective probability model. This objective function is more aligned with the preferences of the consumers than the paternalistic objective function of the benevolent government.

In studying the equilibrium consequences of this change, I follow the same steps as above. I begin this section by formulating both the sequential and recursive versions of the planner's problem. Then, I compute the numerical solutions of this model, comparing the equilibrium to that in a full confidence setup.

Definition 3 *The Ramsey problem of the political government is to choose the competitive equilibrium that maximizes the expected utility of the consumers under the consumers' subjective probability model. The Ramsey outcome under the political government is the competitive equilibrium that attains the maximum.*

Proposition 2 *The allocation and distortions in a Ramsey outcome under the political government solve the following problem:*

$$\max_{c_t, x_t, V_t, b_{t+1}, m_{t+1}, M_{t+1}} \sum_{t=0}^{\infty} \sum_{g^t} \beta^t \pi\left(g^t\right) M_t u\left(c_t, x_t\right)$$

subject to

$$\sum_{t=0}^{\infty}\sum_{g^t}\beta^t\pi\left(g^t\right)M_t\left[u_c\left(c_t,x_t\right)c_t-u_x\left(c_t,x_t\right)\left(1-x_t\right)\right]=u_c\left(c_0,x_0\right)b_0$$

$$m_{t+1}=\frac{\exp\left(\frac{-V_{t+1}}{\theta}\right)}{\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^t\right)\exp\left(\frac{-V_{t+1}}{\theta}\right)}$$

$$V_t=u\left(c_t,x_t\right)+\beta\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^t\right)\left\{m_{t+1}V_{t+1}+\theta m_{t+1}\ln m_{t+1}\right\}$$

$$M_{t+1}=m_{t+1}M_t$$

$$c_t+x_t+g_t=1$$

The proof is similar to the previous one – except one must sum across t to derive the first implementability constraint – and so is suppressed. The transversality condition imposed here is

$$lim_{T\to\infty}\beta^T M_T\lambda_T p_T b_T=0$$

It is clear from this proposition that the government places the same weights on future events as the representative consumer does, even though the government does not share the consumers' lack of confidence. The first constraint is the infinite-time implementability constraint. This constraint differs from the rational expectations version because the distortion M_t affects the perceived probability of each history. In addition, there is one constraint that must be applied to the political government's problem that was not applied to the benevolent government's problem. This constraint tracks the movement in M_t.

3.2.1 Sequential Formulation of the Political Planner's Problem:

The objective function of the political government leads to the following sequential problem, in which the government chooses paths for $\left(c_t,x_t,V_t,\{m_{t+1},M_{t+1}\}\right)$:

$$\mathcal{L}=\sum_{t=0}^{\infty}\beta^t\sum_{g^t}\pi\left(g^t\right)\left\{\begin{array}{c}M_t u\left(c_t,x_t\right)+\mu_t\left[c_t+x_t+g_t-1\right]\\+\xi M_t\left[u_c\left(c_t,x_t\right)c_t-u_x\left(c_t,x_t\right)\left(1-x_t\right)\right]-\xi u_c\left(c_0,x_0\right)b_0\\+M_t\Gamma_t\left[V_t-u\left(c_t,x_t\right)-\beta\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^t\right)\left\{m_{t+1}V_{t+1}+\theta m_{t+1}\ln m_{t+1}\right\}\right]\\+\beta M_t\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^t\right)\varpi_{t+1}\left[m_{t+1}-\frac{\exp\left(\frac{-V_{t+1}}{\theta}\right)}{\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^t\right)\exp\left(\frac{-V_{t+1}}{\theta}\right)}\right]\\+\beta\sum_{g_{t+1}}\pi\left(g_{t+1}\mid g^t\right)\chi_{t+1}\left[M_{t+1}-m_{t+1}M_t\right]\end{array}\right\}$$

The first order conditions for this problem are

$$\begin{aligned} c_t, \forall t &\geq 1: \\ 0 &= u_c(c_t, x_t) + \frac{\mu_t}{M_t} - \Gamma_t u_c(c_t, x_t) \\ &\quad + \xi \left[u_{cc}(c_t, x_t) c_t + u_c(c_t, x_t) - u_{cx}(c_t, x_t)(1 - x_t) \right] \end{aligned}$$

$$\begin{aligned} x_t\left(g^t\right), \forall t &\geq 1: \\ 0 &= u_x(c_t, x_t) + \frac{\mu_t}{M_t} - \Gamma_t u_x(c_t, x_t) \\ &\quad + \xi \left[u_{cx}(c_t, x_t) c_t + u_x(c_t, x_t) - u_{xx}(c_t, x_t)(1 - x_t) \right] \end{aligned}$$

$$\begin{aligned} m_{t+1}, \forall t &\geq 0: \\ 0 &= -\Gamma_t \left[V_{t+1} + \theta (1 + \log m_{t+1}) \right] + \varpi_{t+1} - \chi_{t+1} \end{aligned}$$

$$\begin{aligned} V_t, \forall t &\geq 1: \\ 0 &= \Gamma_t - \Gamma_{t-1} + \left(\frac{1}{\theta} \right) \left[\varpi_t - \sum_{g_t} \pi\left(g_t \mid g^{t-1}\right) m_t \varpi_t \right] \end{aligned}$$

$$\begin{aligned} M_{t+1}, \forall t &\geq 0: \\ 0 &= \chi_{t+1} + u(c_{t+1}, x_{t+1}) + \xi \left[u_c(c_{t+1}, x_{t+1}) c_{t+1} - u_x(c_{t+1}, x_{t+1})(1 - x_{t+1}) \right] \\ &\quad - \beta \sum_{g_{t+2}} \pi\left(g_{t+2} \mid g^{t+1}\right) m_{t+2} \chi_{t+2} \end{aligned}$$

Just as in the benevolent government's first order conditions, Γ_t is a martingale, this time with respect to the subjective expectation. This implies that Γ_{t-1} is a state variable in the recursive version of the problem. Interestingly, the level of the probability distortion, M_t, disappears from the first order conditions. This is because the planner does not seek to re-align the consumers' expectation with the approximating model, since the two agents optimize with respect to the same probability model. As a result, the planner does not need a third state variable to track the level of the distortion. Thus, the corresponding recursive problem only has two state variables.

3.2.2 Recursive Formulation of the Political Planner's Problem:

The planner's problem in recursive form is

$$W(\Gamma_-, g_-) = \min_{\Gamma_g} \max_{c_g, x_g, m_g, V_g} \sum_g \pi(g \mid g_-) \left\{ \begin{array}{c} m_g u(c_g, x_g) + \mu_g [c_g + x_g + g - 1] \\ +\xi m_g [u_c(c_g, x_g) c_g - u_x(c_g, x_g)(1 - x_g)] \\ -\Gamma_- [m_g V_g + \theta m_g \ln m_g] + m_g \Gamma_g [V_g - u(c_g, x_g)] \\ +\varpi_g \left[m_g - \frac{\exp\left(\frac{-V_g}{\theta}\right)}{\sum_g \pi(g|g_-) \exp\left(\frac{-V_g}{\theta}\right)} \right] + \beta m_g W(\Gamma_g, g) \end{array} \right\}$$

The first order and envelope conditions from this problem are described in Appendix B.

The solution to this problem is indexed by the multiplier ξ. For each value of ξ, the first order conditions imply an optimal allocation. In order to solve for the correct value of ξ, the one that corresponds to the initial level of debt held by the consumers, I will find the value such that the implementability constraint is satisfied with equality. For this ξ, the allocation satisfies all constraints and yields the highest subjective welfare for the consumers.

3.2.3 Model Solution and Discussion:

The solution to the political government's problem has been calculated using the same utility function, government spending process, and parameter values as for the benevolent government problem.

Given that the preferences of the two agents are aligned, the intuition underlying the chosen fiscal policy is relatively straight-forward. To illustrate the logic behind the solution, I first go through a small thought experiment. Upon completion, I give a more detailed description that characterizes how the choice of fiscal policy shapes the economy's path and maximizes the consumers' subjective expected utility.

In Lucas and Stokey (1983), it is shown that public debt should be used as insurance against the spending shock. This insurance allows the government to reduce the fluctuations in the labor tax rate, decreasing the volatility of the intra-temporal distortion. Because of the concavity of the consumers' utility function, a smooth intra-temporal distortion comes with the smallest welfare cost for the consumers.

The political government faces an additional reason to limit the volatility in the tax rate: a smooth labor tax reduces the consumers' probability distortion. This, in turn, has the direct effect of increasing the consumers' subjective expected utility.

To see this, first consider the opposite situation. Consider a tax profile that results in a volatile allocation and welfare profile across states. The consumers, in their uncertainty, respond to this volatility by decreasing the weight they place on the high welfare state and increasing the weight they place on the low welfare state. This implies that the consumers' subjective welfare is relatively low, since it is a weighted (weighted by the consumers' subjective probability model) average of the states' welfare.

Now, consider the alternative: the government changes its fiscal policy by lowering the labor tax in war and raising the tax in peace. This tax profile leads to a smoother allocation for the consumers, reducing the volatility of their subjective welfare across state. As a result, the consumers do not fear the low welfare state as much because reaching that state is no longer as harmful. This means that, for a given level of model uncertainty, the consumers reduce their incremental probability distortion. This leads to a larger subjective welfare for the consumers.

This intuition can be seen graphically below. As before, I have plotted the equilibrium solutions for different levels of consumer uncertainty. A large θ implies that the consumers are confident in their approximating model. As θ falls, the consumers are increasingly uncertain about the true model and so raise their probability distortion.

The logic above suggests that the planner decreases (increases) the labor tax in war (in peace), relative to when consumers face no uncertainty. Although the tax rate changes are partially offset by the change in labor supply, the new policy yields a less volatile profile of tax revenues across states. This can be seen in Figure 6. One direct implication of this tax revenue profile is that the volatility of the primary deficit has increased relative to when consumers are confident about the probability model. This means that the government borrows more from consumers to finance its high spending and saves more during low spending states. Further, the more uncertainty the consumers face, the greater the volatility of the primary deficit.

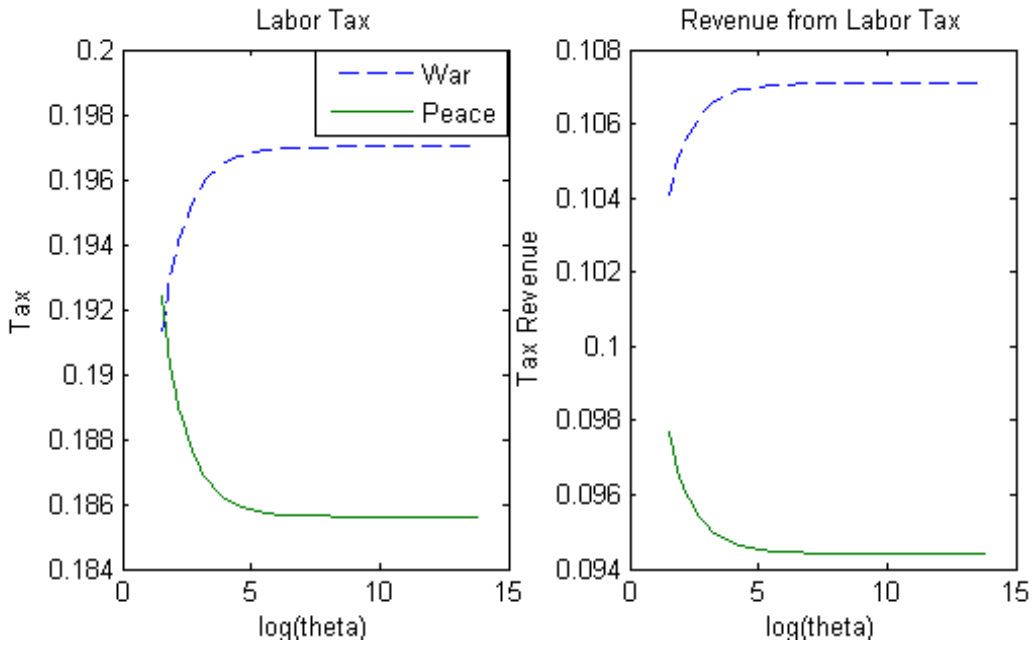

Figure 6: Optimal fiscal policy implemented by the political government for different levels of θ

Critically, this choice of fiscal policy is meant to reduce the volatility of the consumers' welfare across states, since this decreases the size of the consumers' probability distortion. As shown in Figure 7, the combination of raising the consumers' consumption and labor supply during war evidently leads to a smaller decrease in welfare than the peace-time decrease in consumption and labor supply.

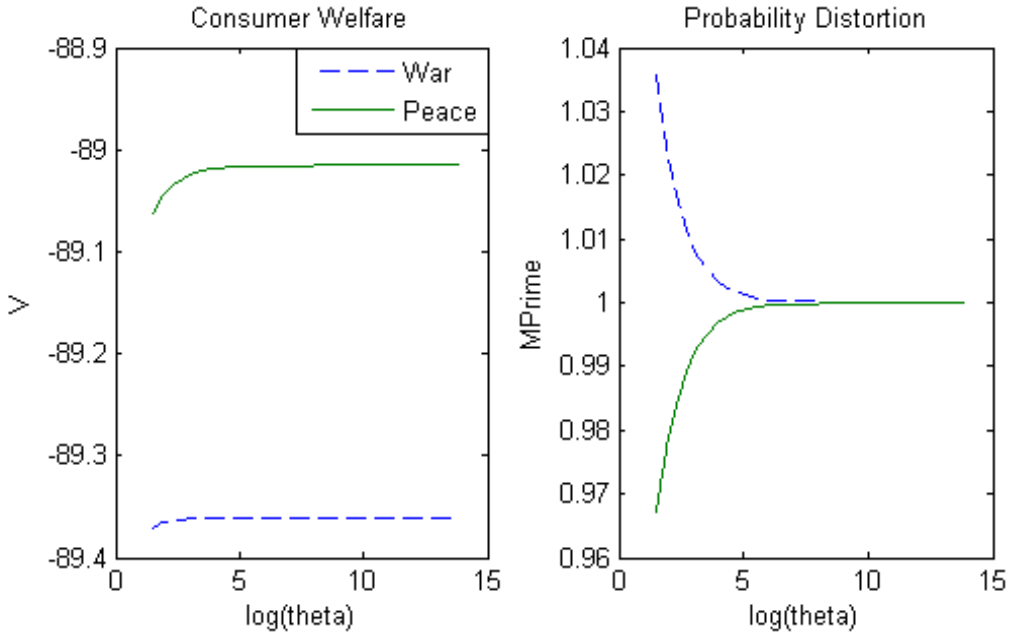

Figure 7: Consumer welfare and probability distortion for different levels of θ

In the graphs below, I plot the impulse response functions to a positive government spending shock under two different degrees of model uncertainty. Again, for all periods except period 1, $g_t = \bar{g}$. In period 1, there is an unexpected increase in government spending so that $g_1 = g_{high}$. The line associated with consumers facing a large degree of model uncertainty is labeled "Uncertain", while the line associated with consumers completely understanding the spending process is labeled "Certain".

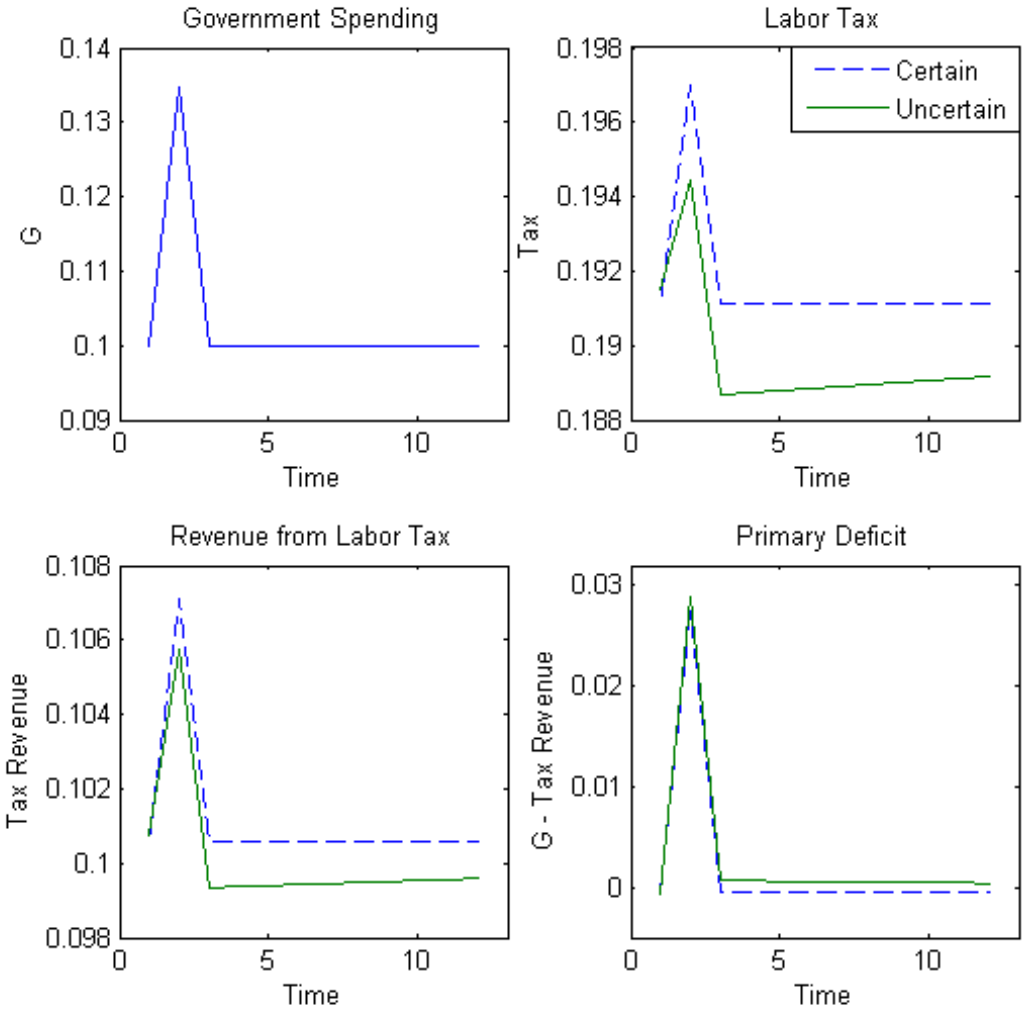

Figure 8: Comparing the policy impulse response functions to an increase in government spending when $\theta \to \infty$ and $\theta = 10$

As discussed above, a positive, one-period war leads the political government to increase its labor tax. This increase, though, is smaller when consumers face uncertainty. Even after the shock has passed, the labor tax is persistently lower when the consumers face model uncertainty. The same results hold for

labor tax revenues, as seen in the bottom left graph of Figure 8. This result suggests that the government must finance a larger portion of the war through contemporaneous borrowings.

4 Model Comparison:

This section briefly compares the fiscal policies implemented by the benevolent and political governments. Both governments seek to use their fiscal instruments in order to minimize two social costs. First, each government wants to minimize the distortion that arises when financing its spending with a linear labor tax. Second, each government wants to minimize the costs associated with consumer uncertainty.

Even though both governments have similar goals, the implied optimal fiscal policy for each government is vastly different. The fundamental reason behind this is a difference in preferences: the expectation used in the political government's objective function is aligned with the consumers' expectation, while the same does not hold for the benevolent government. This alignment is critical because it determines whether or not the government wants to 'correct' the behavior of the consumers. If the government does not seek to re-align the consumers' behavior with what is optimal when the consumers face no uncertainty, then the task of the government is straight-forward: minimize the fluctuations in the consumers' subjective welfare across states. This requires a smooth profile of labor taxes across states. Once accomplished, the consumers' probability distortion shrinks because the consumers are relatively indifferent about which state occurs. The reduced distortion increases the subjective welfare of the consumers.

Figure 9 plots the variance of the consumers' value function across both types of government. As can be seen, the political government smoothes the consumers' welfare across states.

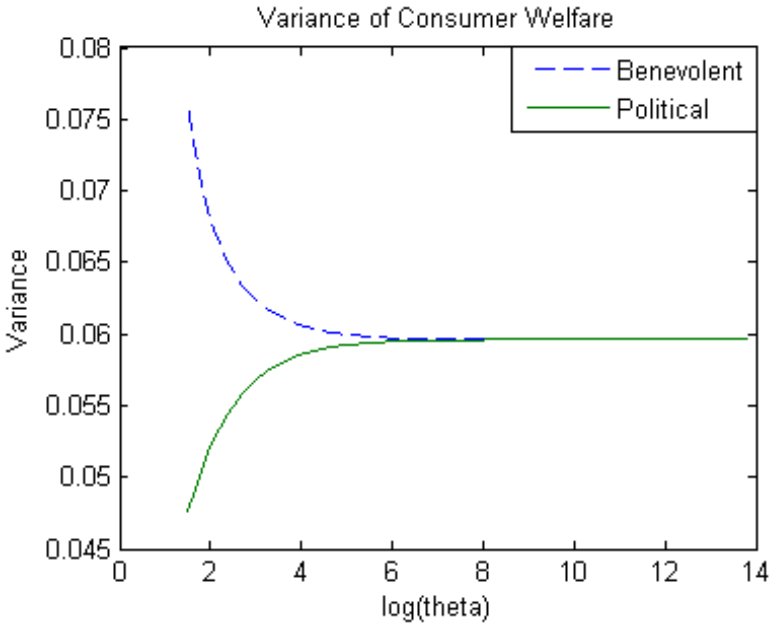

Figure 9: Comparing the consumers' subjective welfare under the benevolent and political governments

The benevolent government, though, responds in the opposite manner. The benevolent government does not attempt to reduce the fluctuations in the consumers' subjective welfare, but instead increases that variance. This action comes with the cost of increasing the probability distortion of the consumers. The benevolent government, though, is willing to accept a larger distortion due to consumer uncertainty because it gains something more valuable in return: the government alters the consumers' expectations so that their savings decisions are now more aligned with the full-confidence setup. This is important to the benevolent government because it optimizes with respect to a different expectation than does the consumer. Thus, the alignment of the expectations is a critical factor in understanding whether the government wants to reduce or increase the consumers' welfare fluctuations across states.

The degree to which the government wants to smooth the consumers' welfare determines, in turn, the optimal correlation between the labor tax and government spending. Because the political government wants to smooth welfare, its labor tax must remain fairly smooth across states and so the correlation remains small. Intuitively, this implies that the labor tax is not used to absorb much of the fiscal shock. Rather, the political government must rely on its borrowings from consumers to finance the shock. Conversely, the benevolent government sets a more volatile labor income tax, raising the tax rate in states of war and lowering it in states of peace. As the correlation between the labor tax rate and government spending is large, the benevolent government does use the tax rate to help absorb the fiscal

shock. Evidence of this can be seen in Figure 10.

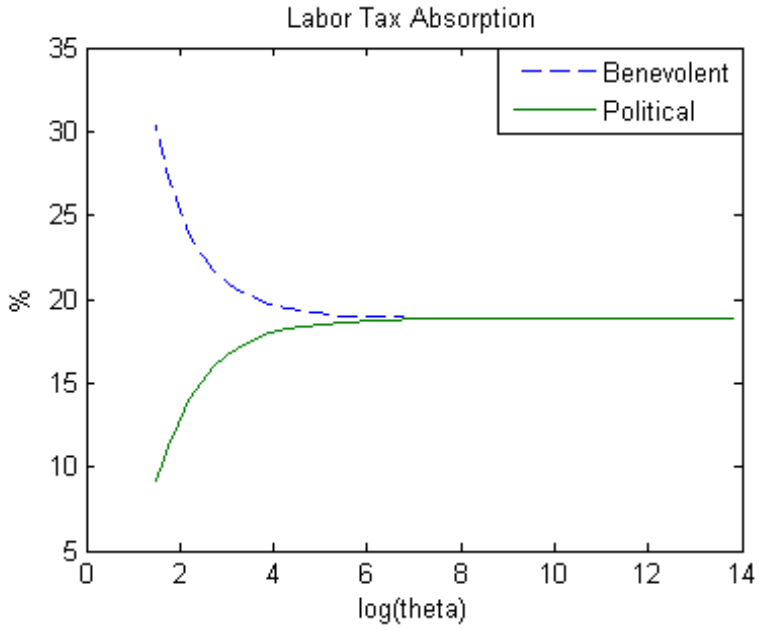

Figure 10: Comparing the degree to which labor taxes finance government spending under the benevolent and political governments

4.0.4 Welfare Comparison:

This section examines the welfare consequences of consumer uncertainty across the benevolent and political governments. One would expect that the benevolent government's policy should achieve a higher welfare for the consumers than the political government's policy. This is because the spending process happens to evolve according to the model used in the benevolent government's problem. The question then is the size of the welfare cost associated with political government relative to the benevolent government. To determine this value, I simulate the path of the economy 100 times. For each simulation, the economy runs for 200 periods. In each of those periods, I calculate the consumers' period utility $u(c_t, x_t)$. Discounting these values back to time 0, I obtain the consumers' lifetime utility. I arrive at the average lifetime welfare by averaging this value across all simulations. I compare this number to W_0, the benevolent government's value of the value function at time 0, which is equal to the consumers' expected utility under the true probability measure. The values are listed in Table 2.

	Benevolent government	Political Government
Welfare	-89.1815	-89.2428

Table 2: Welfare comparisons under the benevolent and political governments

As predicted, the welfare of the consumers is lower under the political government. Following the Lucas treatment, the percentage increase in consumption that must be given to the consumers living under the political government to make them indifferent to living under the benevolent government is 0.02% per period.

5 Conclusion:

This paper compares the fiscal policies chosen by two types of altruistic government when confronted with consumer uncertainty. The first government, labeled a benevolent government, maximizes the consumers' expected utility under the approximating probability model. The second government, labeled a political government, maximizes the consumers' expected utility under the consumers' subjective probability model. Even though both governments face the same distortions in the economy, each responds in a dramatically different way. The benevolent government increases the volatility of the labor tax, which smoothes the fluctuations in the primary deficit across states. The political government, conversely, decreases the volatility of the labor tax, which amplifies the fluctuations in the primary deficit.

Fundamentally, these stark policy differences hinge upon whether the probability model used by the government is aligned with the subjective probability model of the consumers. In the case of the benevolent government, the two measures do not coincide. This leads the benevolent government to choose a policy that 'corrects' the behavior of the consumers. That is, the planner attempts to re-align the decisions of the consumers with their full-confidence values. The most critical decision that the benevolent government wants to re-align is the consumers' savings decision. To do this, the planner sets policy to raise the price of the war-contingent bond and lower the price of the peace-contingent bond. These movements influence the returns on these bonds, reducing the degree to which the consumers' savings profile is different from its full-confidence value.

In the case of the political government, the two probability measures coincide. Because of this, the political government does not attempt to correct the consumers' behavior. Rather, the government chooses policy that reduces the consumers' probability distortion. This is beneficial because it prevents the consumers from hurting their subjective welfare due to their worries about the future.

In this model, I have assumed that the actual stochastic process for the government spending shock happens to be the same as the approximating probability model. As such, the government's confidence in

its probability model is well-placed. This assumption also implies that the policy chosen by the benevolent government leads to higher consumer welfare than the policy chosen by the political government. This need not be the case, though. If the actual process does not correspond to the approximating model, then the government's confidence would be misplaced. In this situation, the political government's policy could lead to higher welfare for the consumers, and the paternalism of the benevolent government would be detrimental.

6 Acknowledgements:

I would like to thank Michael Woodford, Marc Giannoni, and Bruce Preston for their invaluable advice. I am deeply indebted to Stefania Albanesi, Alessandra Casella, John Donaldson, Pierre Yared and the Columbia Seminar participants. All remaining errors are my own.

7 Appendix A:

This appendix describes the first order and envelope conditions from the benevolent planner's recursive problem.

The first order conditions are

$$\begin{aligned} c_g \quad : \quad & 0 = u_c\left(c_g, x_g\right) + \mu_g + \xi_- u_{cc}\left(c_g, x_g\right) m_g b_g - \Gamma_g u_c\left(c_g, x_g\right) \\ & + \xi_g \left[u_{cc}\left(c_g, x_g\right)\left(c_g - b_g\right) + u_c\left(c_g, x_g\right) - u_{cx}\left(c_g, x_g\right)\left(1 - x_g\right)\right] \end{aligned} \tag{20}$$

$$\begin{aligned} x_g \quad : \quad & 0 = u_x\left(c_g, x_g\right) + \mu_g + \xi_- u_{cx}\left(c_g, x_g\right) m_g b_g - \Gamma_g u_x\left(c_g, x_g\right) \\ & + \xi_g \left[u_{cx}\left(c_g, x_g\right)\left(c_g - b_g\right) + u_x\left(c_g, x_g\right) - u_{xx}\left(c_g, x_g\right)\left(1 - x_g\right)\right] \end{aligned} \tag{21}$$

$$m_g : 0 = \xi_- u_c\left(c_g, x_g\right) b_g - \Gamma_- \left[V_g + \theta\left(1 + \ln m_g\right)\right] + \varpi_g \tag{22}$$

$$V_g : 0 = -\Gamma_- m_g + \Gamma_g + \left(\frac{1}{\theta}\right) m_g \left[\varpi_g - \sum_g \pi\left(g \mid g_-\right) m_g \varpi_g\right] \tag{23}$$

$$b_g : 0 = \xi_- m_g - \xi_g \tag{24}$$

$$\xi_g : 0 = u_c\left(c_g, x_g\right)\left(c_g - b_g\right) - u_x\left(c_g, x_g\right)\left(1 - x_g\right) + \beta W_\xi\left(\xi_g, \Gamma_g, g\right) \tag{25}$$

$$\Gamma_g : 0 = V_g - u\left(c_g, x_g\right) + \beta W_\Gamma\left(\xi_g, \Gamma_g, g\right) \tag{26}$$

and the envelope conditions are

$$W_{\xi}\left(\xi_{-},\Gamma_{-},g_{-}\right)=\sum_{g}\pi\left(g\mid g_{-}\right)\left[u_{c}\left(c_{g},x_{g}\right)m_{g}b_{g}\right]$$

$$W_{\Gamma}\left(\xi_{-},\Gamma_{-},g_{-}\right)=-\sum_{g}\pi\left(g\mid g_{-}\right)\left[m_{g}V_{g}+\theta m_{g}\ln m_{g}\right]$$

The envelope conditions can be combined with (25) and (26) to retrieve the implementability constraints. It can be shown that the other four first order conditions are equivalent to those in the sequential formulation. As in the sequential version, the co-state variables are martingales.

The time 0 value function is

$$W_{0}=\min_{\xi_{0},\Gamma_{0}}\max_{c_{0},x_{0},V_{0}}\left\{\begin{array}{c}u\left(c_{0},x_{0}\right)+\mu_{0}\left[c_{0}+x_{0}+g_{0}-1\right]\\+\xi_{0}\left[u_{c}\left(c_{0},x_{0}\right)\left(c_{0}-b_{0}\right)-u_{x}\left(c_{0},x_{0}\right)\left(1-x_{0}\right)\right]\\+\Gamma_{0}\left[V_{0}-u\left(c_{0},x_{0}\right)\right]+\beta W\left(\xi_{0},\Gamma_{0},g_{0}\right)\end{array}\right\}$$

The first order conditions are

$$c_{0}:0=u_{c}\left(c_{0},x_{0}\right)+\mu_{0}+\xi_{0}\left[u_{cc}\left(c_{0},x_{0}\right)\left(c_{0}-b_{0}\right)+u_{c}\left(c_{0},x_{0}\right)-u_{cx}\left(c_{0},x_{0}\right)\left(1-x_{0}\right)\right]-\Gamma_{0}u_{c}\left(c_{0},x_{0}\right)$$

$$x_{0}:0=u_{x}\left(c_{0},x_{0}\right)+\mu_{0}+\xi_{0}\left[u_{cx}\left(c_{0},x_{0}\right)\left(c_{0}-b_{0}\right)+u_{x}\left(c_{0},x_{0}\right)-u_{xx}\left(c_{0},x_{0}\right)\left(1-x_{0}\right)\right]-\Gamma_{0}u_{x}\left(c_{0},x_{0}\right)$$

$$V_{0}:0=\Gamma_{0}$$

$$\xi_{0}:0=u_{c}\left(c_{0},x_{0}\right)\left(c_{0}-b_{0}\right)-u_{x}\left(c_{0},x_{0}\right)\left(1-x_{0}\right)+\beta W_{\xi}\left(\xi_{0},\Gamma_{0},g_{0}\right)$$

$$\Gamma_{0}:0=V_{0}-u\left(c_{0},x_{0}\right)+\beta W_{\Gamma}\left(\xi_{0},\Gamma_{0},g_{0}\right)$$

8 Appendix B:

This appendix describes the first order and envelope conditions from the political planner's recursive problem.

The first order conditions are

$$\begin{aligned} c_g \quad : \quad & 0 = u_c\left(c_g, x_g\right) - \Gamma_g u_c\left(c_g, x_g\right) + \frac{\mu_g}{m_g} \\ & + \xi\left[u_{cc}\left(c_g, x_g\right) c_g + u_c\left(c_g, x_g\right) - u_{cx}\left(c_g, x_g\right)\left(1 - x_g\right)\right] \end{aligned}$$

$$\begin{aligned} x_g \quad : \quad & 0 = u_x\left(c_g, x_g\right) - \Gamma_g u_x\left(c_g, x_g\right) + \frac{\mu_g}{m_g} \\ & + \xi\left[u_{cx}\left(c_g, x_g\right) c_g + u_x\left(c_g, x_g\right) - u_{xx}\left(c_g, x_g\right)\left(1 - x_g\right)\right] \end{aligned}$$

$$\begin{aligned} m_g \quad : \quad & 0 = u\left(c_g, x_g\right) + \xi\left[u_c\left(c_g, x_g\right) c_g - u_x\left(c_g, x_g\right)\left(1 - x_g\right)\right] \\ & - \Gamma_-\left[V_g + \theta\left(1 + \ln m_g\right)\right] + \Gamma_g\left[V_g - u\left(c_g, x_g\right)\right] + \varpi_g + \beta W\left(\Gamma_g, g\right) \end{aligned}$$

$$V_g : 0 = \Gamma_g - \Gamma_- + \left(\frac{1}{\theta}\right)\left[\varpi_g - \sum_g \pi\left(g \mid g_-\right) m_g \varpi_g\right]$$

$$\Gamma_g : 0 = V_g - u\left(c_g, x_g\right) + \beta W_\Gamma\left(\Gamma_g, g\right)$$

where the envelope condition is

$$W_\Gamma\left(\Gamma_-, g_-\right) = -\sum_g \pi\left(g \mid g_-\right)\left[m_g V_g + \theta m_g \ln m_g\right]$$

These first order conditions are equivalent to those derived in the sequential formulation.

The time 0 planner's problem is

$$W_0 = \min_{\Gamma_0} \max_{c_0, x_0, V_0} \left\{ \begin{array}{c} u\left(c_0, x_0\right) + \mu_0\left[c_0 + x_0 + g_0 - 1\right] \\ + \xi\left[u_c\left(c_0, x_0\right)\left(c_0 - b_0\right) - u_x\left(c_0, x_0\right)\left(1 - x_0\right)\right] \\ + \Gamma_0\left[V_0 - u\left(c_0, x_0\right)\right] + \beta W\left(\Gamma_0, g_0\right) \end{array} \right\}$$

and the associated first order conditions are

$$c_0 : 0 = u_c\left(c_0, x_0\right) + \mu_0 + \xi\left[u_{cc}\left(c_0, x_0\right)\left(c_0 - b_0\right) + u_c\left(c_0, x_0\right) - u_{cx}\left(c_0, x_0\right)\left(1 - x_0\right)\right] - \Gamma_0 u_c\left(c_0, x_0\right)$$

$$x_0 : 0 = u_x\left(c_0, x_0\right) + \mu_0 + \xi\left[u_{cx}\left(c_0, x_0\right)\left(c_0 - b_0\right) + u_x\left(c_0, x_0\right) - u_{xx}\left(c_0, x_0\right)\left(1 - x_0\right)\right] - \Gamma_0 u_x\left(c_0, x_0\right)$$

$$V_0 : 0 = \Gamma_0$$

$$\Gamma_0 : 0 = V_0 - u\left(c_0, x_0\right) + \beta W_\Gamma\left(\Gamma_0, g_0\right)$$

References

[1] Camerer, Colin and Roberto Weber. 1992. "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity." *Journal of Risk and Uncertainty.* V, 325-370.

[2] Dennis, Richard. 2007. "Model Uncertainty and Monetary Policy." San Francisco Federal Reserve Bank working paper.

[3] Ellsberg, Daniel. 1961. "Risk, Ambiguity and the Savage Axioms." *Quarterly Journal of Economics.* LXXV, 643-669.

[4] Hansen, Lars and Thomas Sargent. 2001. "Robust Control and Model Uncertainty." *American Economics Review.* 91: 20-66.

[5] Hansen, Lars and Thomas Sargent. 2005. "Robust Estimation and Control Under Commitment." *Journal of Economic Theory.* 124: 258-301.

[6] Hansen, Lars and Thomas Sargent. 2007. "Recursive Robust Estimation and Control Without Commitment." *Journal of Economic Theory.* 136: 1-27.

[7] Karantounias, Anastasios, Lars Hansen, and Thomas Sargent. 2009. "Managing Expectations and Fiscal Policy." Federal Reserve Bank of Atlanta working paper 2009-29.

[8] Kocherlakota, Narayana and Christopher Phelan. 2009. "On the Robustness of Laissez-Faire." *Journal of Economic Theory.* 144: 2372-2387.

[9] Lucas, Robert and Nancy Stokey. 1983. "Optimal Fiscal and Monetary Policy in an Economy Without Capital. *Journal of Monetary Economics.* 12: 55-93.

[10] Marcet, Albert and Ramon Marimon. 1998. "Recursive Contracts." unpublished manuscript.

[11] Orphanides, Athanasios and John Williams. 2007. "Robust Monetary Policy with Imperfect Knowledge." *Journal of Monetary Economics.* 54 (5), 1406-1435.

[12] Svec, Justin. 2010. "Optimal Capital Taxation and Consumer Uncertainty." unpublished manuscript.

[13] Woodford, Michael. 2010. "Robustly Optimal Monetary Policy with Near-Rational Expectations." *American Economic Review.* 100: 274-303.