Optimal Capital Taxation and Consumer Uncertainty

Ryan Chahrour*
Boston College
Toulouse School of Economics

Justin Svec†
College of the Holy Cross

March 31, 2014

Abstract

This paper analyzes the impact of consumer uncertainty on optimal fiscal policy in a model with capital. The consumers lack confidence about the probability model that characterizes the stochastic environment and so apply a max-min operator to their optimization problem. An altruistic fiscal authority does not face this Knightian uncertainty. We show analytically that, in responding to consumer uncertainty, the government no longer sets the expected capital tax rate exactly equal to zero, as is the case in the full-confidence benchmark model. Rather, our numerical results indicate that the government chooses to subsidize capital income, albeit at a modest rate. We also show that the government responds to consumer uncertainty by smoothing the labor tax across states and by making the labor tax persistent.

JEL Codes: E62, H21

Keywords: Model uncertainty, capital income tax, public debt

*Department of Economics, Boston College, Chestnut Hill, MA, U.S.A. and Toulouse School of Economics, Toulouse, France. email: ryan.chahrour@bc.edu.

†Department of Economics, College of the Holy Cross, Worcester MA, USA. email: jsvec@holycross.edu

1 Introduction:

In the typical public finance model with rational expectations, fiscal policy can influence consumer behavior by manipulating the consumers' expectations. That is, by committing to future policy, the government shapes the consumers' beliefs about the possible paths of the endogenous variables, which then affects the consumers' behavior in earlier periods. The assumption of rational expectations helps facilitate this pathway, enabling the consumers to correctly forecast both the state-contingent values of the endogenous variables and the probability model over these variables.

Rational expectations, though, might exaggerate the ability of consumers to understand the stochastic equilibrium. This exaggeration could be costly in that it might mean that the typical fiscal policy model overemphasizes how precisely consumers respond to future policy commitments of the government. If instead consumers face uncertainty about the economy's true probability model, their expectations and behavior might be quite different than those predicted in a rational expectations model. As a consequence, the fiscal authority might find it optimal to implement a different set of fiscal policies knowing that the consumers face model uncertainty.

Karantounias (2013) and Svec (2011) are two examples that introduce consumer uncertainty in an optimal fiscal policy model. In these models without capital, the authors show that the consumers' uncertainty does indeed alter the government's policy decisions. This is because fiscal policy must mitigate the welfare costs associated with both linear taxes and consumer uncertainty. Depending on the specific type of altruism exhibited by the planner, the optimal policy involves either more or less reliance on the labor income tax to finance public spending than is optimal under the baseline model in which consumers do not face model uncertainty.

Although these results are suggestive, the impact of consumer uncertainty on optimal fiscal policy should be most salient in a model with capital, as the consumers' expectations are of primary importance in the design of optimal policy. To this end, this paper introduces consumer uncertainty into the neoclassical growth model of Chari, Christiano, and Kehoe (1994). We formalize the consumers' uncertainty and their resulting behavior by following Hansen and Sargent (2001, 2005, 2007) and the robust control literature. In this approach, consumers are unsure which probability model characterizes the random shocks to government expenditure. They believe that the true probability model lies somewhere within a range of alternative probability models. Each alternative model is represented as a martingale perturbing the approximating probability model. With this type of uncertainty, the robust control literature assumes that the consumers optimize according to max-min preferences, choosing the allocation that maximizes their expected utility, where the expectation is taken with respect to the probability model that minimizes their welfare. The resulting allocation is labeled the robustly optimal allocation, and the worst-case

probability model is labeled the consumers' subjective probability model. This behavior helps ensure that the consumers' utility never falls too far, regardless of which probability model happens to be correct.

Although it is assumed that the consumers are uncertain as to the correct probability model, the opposite assumption is made for the fiscal authority: the government is fully confident that the approximating probability model truly characterizes the stochastic environment. This confidence might be due to the fact that the economy's only source of randomness is a shock to the government's own spending, a process that the government supposes it knows well. To be clear, the consumers and the government are both endowed with the same approximating model, a model that specifies the probability model associated with the exogenous and endogenous variables. But, only the government trusts that this approximating model correctly describes the economy's randomness. The consumers, on the other hand, doubt the accuracy of this model and, perhaps due to a lack of trust in political institutions, cannot be convinced by the government about its accuracy.[1,2]

Critically, this confidence dichotomy reveals a number of possible objective functions for an altruistic government. These objective functions differ as to which expectation they use to calculate the consumers' expected utility. That is, the government could optimize with respect to the approximating probability model or it could optimize according to any one of the alternative probability models that the consumers believe could describe the economy, including the subjective probability model. As the consumers distrust the government's confidence in the approximating probability model, it is not clear which model an altruistic government should use in its optimization problem.

Given this multiplicity of possible objective functions, the assumption made in this paper is that the fiscal authority maximizes the consumers' expected utility under the consumers' own subjective expectation. This choice can be justified for political economy reasons: because the consumers would ex-ante prefer a government that optimizes according to the same probability model that they use, any government chosen by the consumers must design its policies to maximize the consumers' expected utility under their subjective probability model. That is, even though the government believes that the approximating model is correct, the consumers do not trust the government's belief. As such, the consumers would choose a government that optimizes according to the consumers' subjective probability model. One additional benefit of this choice of objective function is that it allows for a one-step deviation from the rational

[1]The fact that the consumers do not trust any announcement made by the government describing what it believes to be the true probability model is particularly relevant because, as discussed in Woodford (2010), the fiscal authority might have the incentive to misrepresent the true probability model in order to manipulate the consumers' behavior. Understanding this, the consumers are skeptical about any government announcement.

[2]While it seems reasonable that the government has more confidence about the stochastic nature of government spending than do the consumers, an interesting alternative would be to assume the government is also uncertain about the shock process. We leave this extension for future work.

expectations framework, since both the consumers and the planner optimize with respect to the same expectation.

Arguably, a similar combination of preferences was on display in the United States during the recent financial crisis. Under this interpretation, Americans faced uncertainty about whether the US government's budget was sustainable. Assuming that Americans are uncertainty-averse, they responded to this uncertainty by fearing that, with large probability, the government's budget was not sustainable. This altered subjective expectation then affected their behavior, leading among other things to investors purchasing assets that hold their value even in times of budget crisis (gold, for example). Continuing with this analogy, the American government understood that the true probability of default was lower than that feared by its citizens (as the government could borrow at historically low interest rates and it could always finance the debt by printing more money). Despite knowing this, however, the government felt pressured by the American people to take actions *as if* the probability of a budget crisis was large. It was perhaps this pressure that led Congressional Republicans to refuse to raise the debt ceiling in 2011, resulting in the sequester.

With this setup, the optimal policy implemented by the fiscal authority involves one period of transition. During that period, the government subsidizes labor with a negative tax on labor income and implements a large tax on capital income, as in Chari, Christiano, and Kehoe (1994). From that period forward, there are three main properties of the time-invariant optimal policies. First, it can be shown analytically that, under one condition, the expected capital tax rate is non-zero even in the case of log utility preferences, breaking the rational expectations result. To derive the magnitude and direction of this deviation from zero, we numerically solve the model. We find numerically that the government chooses to subsidize the consumers' capital income, on average, at a modest rate.

Second, consumer uncertainty leads the government to choose an even more smooth profile of labor taxes across states than is optimal under the full-confidence benchmark. Specifically, greater uncertainty leads the government to reduce the labor tax during periods of high fiscal expenditure, and conversely to raise the labor tax during periods of low expenditure. As labor taxes absorb even less of the fiscal shock than is optimal under the benchmark, the government raises the state-contingent volatility of the private assets tax, a combination of the return on public debt and capital taxes. Third, the government increases the persistence of the labor tax across time. This last result represents an additional qualitative departure from the benchmark model of Chari, Christiano, and Kehoe (1994), where the labor tax inherits the time-series properties of the underlying shock process.

These deviations from the canonical prescriptions of optimal policy are driven by two distinct motivations faced by the planner, both of which are novel to the consumer uncertainty framework. First, the planner seeks to exploit the consumers' subjective expectation so as to reduce the cost of its fiscal insur-

ance. Second, the planner seeks to smooth the consumers' welfare across states, in order to mitigate the direct welfare costs created by the consumers tilting their subjective probabilities in a pessimistic fashion. As we will show below, both of these motivations imply the same policy response: smooth the labor tax across states.

The current paper fits into a larger strand of the recent literature that analyzes how model uncertainty alters the policy conclusions derived from rational expectations models. Generally, this literature has focused on planner uncertainty within a monetary policy framework; examples include Dennis (2010), Dennis, Leitemo, and Soderstrom (2009), Hansen and Sargent (2008), Leitemo and Soderstrom (2008), Levin and Williams (2003), Onatski and Stock (2002), and Walsh (2004). Woodford (2010) modifies the type of uncertainty considered by assuming that the central bank is uncertain of the expectations held by firms, but not uncertain about the stochastic environment. In addition to examining fiscal policy rather than monetary policy, the current analysis differs from most of the literature by examining the policy implications of consumer uncertainty rather than the planner's uncertainty. Finally, this paper is novel in that, to the best of our knowledge, it is the first to analyze optimal capital income tax rates in a model with consumer uncertainty.

This model also wades into the large literature on the optimal size of the capital income tax. Early works, including Judd (1985) and Chamley (1986), suggest that a fiscal planner should not tax capital income in the long run. As mentioned above, Chari, Christiano, and Kehoe (1994) shows that, under restrictive preference assumptions, this result extends to all periods after an initial period within a neo-classical, stochastic growth model. Further, even after relaxing those strict preference assumptions, the ex-ante capital tax rate remains small.

More recent work, though, has identified a number of conditions under which the optimal tax on capital income is non-zero. Aiyagari (1995) and Imrohoroglu (1998), for example, show that imperfections in financial markets could lead the government to choose a positive capital tax. Correia (1996) shows that when the government cannot tax all factors of production at their optimal levels, then the capital tax could be positive or negative, depending on the technological characteristics of that factor. In a similar vein, Piketty and Saez (2012) and Stockman (2001) show that if the government does not have access to the correct set of policy instruments or is constrained in how it uses those instruments, the government could set a non-zero capital tax. Guo and Lansing (1999) analyzes the impact of imperfect competition in the factor markets on the optimal capital tax and shows that the sign of the capital tax depends on whether the reduction in investment relative to the socially efficient benchmark dominates the ability of the government to tax the economic rents of the firms. Klein and Rios-Rull (2003) shows that the optimal capital tax is positive if the planner is unable to fully commit to its fiscal policy. Further, overlapping generations models often lead to a positive capital tax, as it allows the government to effectively condition

the tax rate on the agent's position in her life-cycle; see Conesa, Kitao, and Krueger (2009) for a fuller description of this result.

The outline of the paper is as follows. Section 2 describes the economic environment and characterizes the type of uncertainty faced by the consumers. The optimization problem of the consumers is also formulated. Section 3 discusses the planner's optimization problem, and derives the analytical result that the fiscal authority no longer sets the ex-ante capital income tax equal to zero. Section 4 describes how we numerically implemented this model, while Section 5 examines the numerical results. Section 6 concludes.

2 The economy:

Time is discrete in this infinite-horizon production economy. There are three types of agents: a government, an infinite number of identical consumers, and firms. The only source of randomness in the model is a shock to government spending. This shock can take on a finite number of values. Let $g^t = (g_0, ..., g_t)$ represent the history of the spending shock up to and including period t, where the probability of each history is $\pi(g^t)$. All variables in the model are contingent on the history of spending shock up to that point. In period 0, government spending is known to be g_0 with probability 1. The government finances this expenditure through either taxes or debt, b_t. The government has access to a labor income tax, τ_t, and a capital income tax, Ω_t. Both are restricted to be proportional taxes. Government debt has a state-contingent return, $R_{b,t}$, and matures in one period. Thus, the government's policy at each g^t is to choose the combination of labor and capital taxes and the return on public debt. The period budget constraint of the government is

$$b_t = R_{b,t} b_{t-1} + g_t - \tau_t w_t l_t - \Omega_t \left[r_t - \delta \right] k_{t-1}. \tag{1}$$

Note that the capital income tax applies to the after-depreciation return on capital, where δ is the depreciation rate.

Each consumer's wealth is composed of three components: after-tax labor income, after-tax capital income, and a return on debt held from the previous period. Out of this wealth, the consumer can choose to consume, buy capital, or save in the debt market. In each period, the consumer also chooses how much labor to supply. The period budget constraint for the consumer is

$$c_t + k_t + b_t \leq (1 - \tau_t) w_t l_t + R_{k,t} k_{t-1} + R_{b,t} b_{t-1}, \tag{2}$$

where $R_{k,t} = 1 + (1 - \Omega_t)(r_t - \delta)$ is the gross, after-tax return on capital.

A constant returns to scale production function, $F(k_{t-1}, l_t)$, transforms labor and capital into output. This production function satisfies the Inada conditions. The resulting output can be used for private

consumption c_t, public consumption g_t, or investment $k_t - (1-\delta)\,k_{t-1}$. The economy-wide resource constraint is therefore

$$c_t + k_t + g_t = F\left(k_{t-1}, l_t\right) + (1-\delta)\,k_{t-1}. \tag{3}$$

Competitive firms ensure that the returns on labor and capital equal their respective marginal products:

$$w_t = F_l\left(k_{t-1}, l_t\right) \tag{4}$$

and

$$r_t = F_k\left(k_{t-1}, l_t\right). \tag{5}$$

2.1 The consumers' model uncertainty:

The consumers are endowed with an approximating probability model that specifies a probability measure over the paths of the exogenous and endogenous variables. Unlike in Chari, Christiano, and Kehoe (1994), the consumers are uncertain whether this approximating model correctly characterizes the equilibrium. Instead, they worry that other probability measures could potentially describe the stochastic nature of the economy. To ensure that these alternative models conform to some degree with the approximating model, restrictions must be placed on what types of alternative models are allowed.

With the assumption of absolute continuity, the Radon-Nikodym Theorem indicates that 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 approximating probability model:

$$\widetilde{E}\left[X_t\right] = E\left[M_t X_t\right] \tag{6}$$

where $E\left[M_t\right] = 1$ and $\widetilde{E}$ is the subjective expectations operator.[3] This is important, as it allows me to recast consumer uncertainty. Earlier, the consumers were described as being uncertain about the probability model that characterizes the paths of the exogenous and endogenous variables; now, the consumers can be viewed as understanding the correct mapping from states of the world to equilibrium outcome, even though they may not place the correct probability on each state.

By defining an additional term, one can begin to measure the distance between an alternative probability model and the approximating probability model. Let the incremental probability distortion be

$$m_{t+1} = \frac{M_{t+1}}{M_t}, \forall M_t > 0 \tag{7}$$

and $m_{t+1} = 1$ otherwise. This incremental distortion must satisfy $E_t m_{t+1} = 1$, implying that the probability distortion M_t is a martingale. This restriction guarantees that the alternative probability

[3] For more information on this process, please see Hansen and Sargent (2005, 2006).

measures are legitimate probability models. With this definition, the one-period distance between the alternative and approximating models is measured by relative entropy:

$$\epsilon_t\left(m_{t+1}\right) \equiv E_t m_{t+1} \log m_{t+1} \tag{8}$$

This measure is convex and grounded, attaining its minimum when $m_{t+1} = 1, \forall g_{t+1}$.

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

$$\Phi_0 = E_0 \sum_{t=0}^{\infty} \beta^t M_t \epsilon_t\left(m_{t+1}\right) \tag{9}$$

This distance measure is used in the multiplier preferences of Hansen and Sargent (2006). The multiplier preferences characterize how the consumers rank their allocations. Given these preferences, the consumers choose the allocation that maximizes the following criteria:

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

where $u\left(c, l\right)$ is increasing in consumption, decreasing in labor, and strictly concave.

Given these preferences, the consumers worry that for any allocation considered, the probability model that turns out to be correct is the alternative model that results in the lowest expected utility for the consumers in the set of possible alternative models. The resulting alternative model, which we call the consumers' subjective expectation, places greater weight on low welfare states of the world and lower weight on high welfare states of the world than does the approximating probability model, as we will show.

The coefficient $\theta > 0$ is a penalty parameter that indexes the degree to which consumers are uncertain about the probability measure. A small θ implies that the consumers are not penalized too harshly for distorting their probability model away from the approximating model. The min operator then yields incremental probability distortions that diverge greatly from one. The resulting probabilities $\{\pi\left(g^t\right) M_t\}$ are distant from the approximating model. Thus, a small θ indicates that consumers are very unsure about the approximating model and so fear a large set of alternative models. A larger θ means that the consumers face a sizable penalty for distorting their probability model away from the approximating model. As a result, the min operator yields incremental distortions close to one, implying that the worst-case alternative model is close to the approximating model. Thus, a large θ signifies that the consumers have more confidence about the underlying measure and so fear only a small set of alternative models. As $\theta \to \infty$, this model collapses to the rational expectations framework of Chari, Christiano, and Kehoe (1994).[4]

[4]Given the economic environment, there is a one-to-one mapping between a value for the total distortion, $\bar{\Phi}$, and the

2.2 The consumer's problem:

With this formalism, the consumer's problem can be written recursively using the value function $V(b_-, k_-, g, A)$:

$$V(b_-, k_-, g, A) = \max_{c,l,b,k} \min_{m'} \left\{ \begin{array}{c} u(c,l) + \beta \sum_{g'} \pi(g' \mid g) \left[m' V(b,k,g',A') + \theta m' \log m' \right] \\ -\lambda \left[c + k + b - (1-\tau) wl - R_b b_- - R_k k_- \right] \\ -\beta\theta\Psi \left[\sum_{g'} \pi(g' \mid g) m' - 1 \right] \end{array} \right\}$$

where A represents the set of aggregate state variables that the consumers must track in order to forecast fiscal policy in all histories. This set of state variables comes from the government's optimization problem. The consumer believes that her decisions cannot affect the movements of these aggregate state variables. In addition to the period budget constraint, the consumer faces the legitimacy constraint, $\sum_{g'} \pi(g' \mid g) m' = 1$, described above.

Solving the consumer's Bellman equation for the robustly optimal allocation is a two-stage process. In the inner minimization stage, the consumer fears that, for a given allocation, the worst-case probability model over the government spending shocks will occur. The solution that results from this minimization is the consumer's subjective expectation. The outer maximization stage determines the allocation that maximizes the consumers' expected utility, taking into account the endogenous tilting of the consumers' expectation. The solution from this stage is the consumer's robustly optimal allocation.

2.2.1 The inner minimization stage:

As indicated above, the minimization stage yields the subjective probability model that minimizes the consumer's expected utility for a given allocation. The resulting subjective model will be a new probability model that places potentially new probabilities over all the possible states of the world. The state-contingent probability distortion, which balances the marginal benefit of lowering the consumer's expected utility with the marginal cost of the convex penalty term, solves the following equation:

$$V(b,k,g',A') + \theta(1 + \log m') - \theta\Psi = 0$$

Combining this first order condition with the legitimacy constraint, the optimal distortion is

$$m' = \frac{\exp\left(\frac{-V(b,k,g',A')}{\theta}\right)}{\sum_{g'} \pi(g' \mid g) \exp\left(\frac{-V(b,k,g',A')}{\theta}\right)} \tag{10}$$

parameter θ, which has the interpretation of a lagrange multiplier on the optimization problem

$$\min_{m_{t+1}, M_{t+1}} \sum_{t=0}^{\infty} \sum_{g^t} \beta^t \pi(g^t) M_t u(c_t, l_t)$$

subject to the constraint $\Phi_0 \leq \bar{\Phi}$. In numerical analysis, we will focus on comparisons for fixed $\bar{\Phi}$.

This equation describes the consumer's worst-case, state-contingent incremental probability distortion. The magnitude and direction of this distortion depend upon the consumer's subjective welfare, V, in each state in period $t+1$. To better understand this function, consider a two-state government spending process. Suppose that the equilibrium allocation yields a high subjective welfare in state A and a low subjective welfare in state B. Plugging these values into (10), we see that $m_A < 1$ and $m_B > 1$. These distortions imply that consumers fear that the likelihood of state A is small and that the likelihood of state B is large relative to the approximating model.

The degree to which these multiplicative distortions diverge from unity depends upon θ and the difference between V_H and V_L. All else equal, a large θ decreases the probability distortion in all states in period $t+1$, meaning that $\{m_{t+1}\}$ remains closer to one. A small θ, conversely, implies that the probability distortions are further away from one. Also, all else equal, as the difference between V_H and V_L grows, the consumer's alternative model is increasingly far from her approximating model.

One final note about equation (10): because the consumers' subjective expectation is endogenous and depends on the state-contingent values of V, the fiscal planner can influence the consumers' expectations through its choice of policy. That is, by choosing policy that raises or lowers the consumers' welfare in a particular state (which could involve adjusting the consumers' utility contemporaneously or in the future), the planner can alter the perceived likelihood that that state occurs. This expectations channel will become important in understanding the motivations behind the chosen fiscal policy.

2.2.2 The outer maximization stage:

In the maximization stage, the consumer chooses the allocation that maximizes her subjective expected utility, where the expectation is taken with respect to the worst-case probability model over government spending. To find this allocation, we have incorporated the subjective probability model that is derived in the minimization stage into the consumer's optimization problem. The resulting Bellman equation is

$$V(b_-, k_-, g, A) = \max_{c,l,b,k} \left\{ \begin{array}{c} u(c,l) - \beta\theta \log \sum_{g'} \pi(g' \mid g) \exp\left(\frac{-V(b,k,g',A')}{\theta}\right) \\ -\lambda\left[c + k + b - (1-\tau) wl - R_b b_- - R_k k_-\right] \end{array} \right\}$$

This equation highlights the fact that the consumer does not weight her future welfare as she would if she were fully confident in the approximating probability model. Rather, the allocation alters the consumer's future subjective welfare, which in turn influences the endogenous probability distortion.

As is standard in fiscal policy models in which the government must set linear taxes, the intra-temporal condition between consumption and labor is

$$-\frac{u_l(c,l)}{u_c(c,l)} = (1-\tau) w \tag{11}$$

This equation links the marginal disutility of labor with the marginal benefit of raising consumption through increased labor supply. The linear labor tax distorts the optimal tradeoff away from the first-best: $-\frac{u_l(c,l)}{u_c(c,l)w} = 1$.

The two inter-temporal conditions are

$$1 = \beta \sum_{g'} \pi\left(g' \mid g\right) m' \frac{u_c\left(c', l'\right)}{u_c\left(c, l\right)} R'_b \tag{12}$$

$$1 = \beta \sum_{g'} \pi\left(g' \mid g\right) m' \frac{u_c\left(c', l'\right)}{u_c\left(c, l\right)} R'_k \tag{13}$$

These equations balance the marginal utility of increasing consumption today with the expected marginal utility from saving that additional unit in the debt or capital markets. Since the consumer faces model uncertainty, the conditional expectation within these equations is taken with respect to the subjective probability model. It is this point that implies that the government now has an additional channel through which it can affect the economy's asset prices: by adjusting its policy, the planner influences the consumers' expectations which, in turn, affects the prices and returns on capital and debt.

The envelope conditions are

$$V_b\left(b_-, k_-, g; A\right) = \lambda R_b$$

$$V_k\left(b_-, k_-, g; A\right) = \lambda R_k$$

Definition 1 *Given an initial allocation* $\{b_{-1}, k_{-1}\}$*, an initial policy value* Ω_0*, and an initial return on debt* $R_{b,0}$*, a competitive equilibrium is a history-dependent allocation* $\{c_t, l_t, b_t, k_t\}_{t=0}^{\infty}$*, probability distortions* $\{m_{t+1}, M_{t+1}\}_{t=0}^{\infty}$*, prices* $\{r_t, w_t\}_{t=0}^{\infty}$*, returns* $\{R_{k,t+1}, R_{b,t+1}\}_{t=0}^{\infty}$*, and fiscal policies* $\{\tau_t, \Omega_t\}_{t=0}^{\infty}$ *such that*

1. *The probability distortion solves the consumer's inner minimization problem*
2. *The allocation solves the consumer's outer maximization problem, and*
3. *The allocation is feasible, satisfying* (3)*.*

3 The government's problem:

This section considers the policy problem of the government. We assume that the government has access to a commitment technology with which it is able to bind itself to a sequence of policies chosen at $t = 0$. Unlike the consumers, the government is fully confident that the approximating probability model accurately describes the government spending process.

As the definition of the competitive equilibrium makes clear, there are a continuum of possible competitive equilibria, each indexed by a fiscal policy $\{\tau_t, \Omega_t\}_{t=0}^{\infty}$. The outcome, then, depends upon the

objective of the fiscal authority. For the purposes of this paper, we assume that the planner maximizes the consumers' expected utility under their subjective probability model. This decision implies that the government optimizes with respect to the same probability model as the consumers. Note, though, that the government does not include the entropy term from the consumers' preferences in its objective function. The result is an objective function that is similar to that of the political government in Svec (2011).

With this choice of planner preferences, the Ramsey outcome is the competitive equilibrium that attains the maximum. In formulating the Ramsey problem, we will follow the primal approach in which the government chooses the consumers' allocation and probability distortions. With these values, we will then back out what fiscal policies implement this competitive equilibrium.

Proposition 1 *The allocation and distortions in a Ramsey outcome solve the following problem:*

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

subject to

$$\sum_{t=0}^{\infty} \sum_{g^t} \beta^t \pi\left(g^t\right) M_t \left[u_c\left(c_t, l_t\right) c_t + u_l\left(c_t, l_t\right) l_t\right] = u_c\left(c_0, l_0\right) \left[R_{b0} b_{-1} + R_{k0} k_{-1}\right] \tag{14}$$

$$m_{t+1} = \frac{\exp\left(\frac{-V_{t+1}}{\theta}\right)}{\sum\limits_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) \exp\left(\frac{-V_{t+1}}{\theta}\right)} \tag{15}$$

$$V_t = u\left(c_t, l_t\right) + \beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) \left\{m_{t+1} V_{t+1} + \theta m_{t+1} \ln m_{t+1}\right\} \tag{16}$$

$$M_{t+1} = m_{t+1} M_t \tag{17}$$

$$c_t + g_t + k_t = F\left(k_{t-1}, l_t, g_t\right) + \left(1 - \delta\right) k_{t-1} \tag{18}$$

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 (14) − (18). To demonstrate this, we first show that any allocation and probability distortion that satisfies the competitive equilibria constraints must also satisfy (14) − (18). Multiply (2) by $\beta^t \pi\left(g^t\right) M\left(g^t\right) \lambda\left(g^t\right)$ and sum over t and g^t. Plugging in (11) − (13) and using the two transversality conditions

$$\begin{aligned} \lim_{T \to \infty} \beta^T M_T \lambda_T b_T &= 0 \\ \lim_{T \to \infty} \beta^T M_T \lambda_T k_T &= 0 \end{aligned}$$

reveals the constraint (14). The constraint (15) follows directly from the optimality condition in the inner minimization, (17) comes from the definition of m_{t+1}, and (16) is the representative consumer's Bellman

equation. Finally, (18) is the resource constraint which ensures feasibility. Thus, $(14)-(18)$ are necessary conditions that the Ramsey outcome must solve. Going in the other direction, given an allocation and distortions that satisfy $(14)-(18)$, policies and prices can be determined from $(1)-(5)$ and the consumer's first order conditions. ■

The first constraint in the planner's problem is the implementability constraint. This constraint differs from its rational expectations counterpart in that the planner must account for the consumers' probability distortion at each date t. This is accomplished by the multiplicative term, M_t. In order to incorporate how policy affects this distortion, the planner must keep track of how that distortion is set and how it is updated across time and state. This information is contained in the next three constraints. The final constraint is the resource constraint.

The proposition above describes the robustly optimal allocation and distortions that achieve the Ramsey outcome. The bond holdings in history g^r that support this competitive equilibrium are described by

$$b_r = \frac{\sum_{t=r+1}^{\infty} \sum_{g^t} \beta^{t-r} \pi\left(g^t \mid g^r\right) M_t \left[u_c\left(c_t, l_t\right) c_t + u_l\left(c_t, l_t\right) l_t\right]}{M_r U c\left(c_r, l_r\right)} - k_r \tag{19}$$

This value is pinned down using the future, state-contingent values of consumption, labor supply, capital, and probability distortions.

As in the previous literature, the government in our economy has the incentive to finance its public spending by raising very large taxes on the inelastic goods of capital and debt at $t = 0$. To prevent this outcome, we assume exogenous values for the initial capital tax, Ω_0, and return on debt $R_{b,0}$.

3.1 Sequential Formulation of Ramsey Problem:

With this setup, we now formulate the government's sequential problem:

$$\begin{aligned}
\mathcal{L} \quad = \quad & \sum_{t=0}^{\infty} \sum_{g^t} \beta^t \pi\left(g^t\right) \left\{ \begin{array}{c}
M_t u\left(c_t, l_t\right) + \xi M_t \left[u_c\left(c_t, l_t\right) c_t + u_l\left(c_t, l_t\right) l_t\right] \\
+ M_t \mu_t \left[c_t + g_t + k_t - F\left(k_{t-1}, l_t, g_t\right) - (1-\delta) k_{t-1}\right] \\
+ M_t \Gamma_t \left[V_t - u\left(c_t, l_t\right) - \beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) \left\{m_{t+1} V_{t+1} + \theta m_{t+1} \ln m_{t+1}\right\}\right] \\
+ \beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) \Phi_{t+1} \left[M_{t+1} - m_{t+1} M_t\right] \\
+ \beta M_t \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) \omega_{t+1} \left[m_{t+1} - \frac{\exp\left(\frac{-V_{t+1}}{\theta}\right)}{\sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) \exp\left(\frac{-V_{t+1}}{\theta}\right)}\right]
\end{array} \right\} \\
& - \xi u_c\left(c_0, l_0\right) \left[R_{b0} b_{-1} + \left\{1 + \left[1 - \Omega_0\right] \left[F_k\left(k_{-1}, l_0, g_0\right) - \delta\right]\right\} k_{-1}\right]
\end{aligned}$$

The first-order necessary conditions for $t \geq 1$ are

$$c_t : u_c\left(c_t, l_t\right) + \xi \left[u_{cc}\left(c_t, l_t\right) c_t + u_c\left(c_t, l_t\right) + u_{cl}\left(c_t, l_t\right) l_t\right] + \mu_t - \Gamma_t u_c\left(c_t, l_t\right) = 0 \tag{20}$$

$$l_t : u_l(c_t, l_t) + \xi\left[u_{cl}(c_t, l_t)\, c_t + u_{ll}(c_t, l_t)\, l_t + u_l(c_t, l_t)\right] - \mu_t F_l(k_{t-1}, l_t, g_t) - \Gamma_t u_l(c_t, l_t) = 0 \tag{21}$$

$$V_t : \Gamma_t - \Gamma_{t-1} + \left(\frac{1}{\theta}\right)\left[\omega_t - \sum_{g_t} \pi\left(g_t \mid g^{t-1}\right) m_t \omega_t\right] = 0 \tag{22}$$

$$k_t : \mu_t - \sum_{g_{t+1}} \beta\pi\left(g_{t+1} \mid g^t\right) m_{t+1}\mu_{t+1}\left[F_k(k_t, l_{t+1}, g_{t+1}) + 1 - \delta\right] = 0 \tag{23}$$

$$M_t : u(c_t, l_t) + \xi\left[u_c(c_t, l_t)\, c_t + u_l(c_t, l_t)\, l_t\right] - \sum_{g_{t+1}} \beta\pi\left(g_{t+1} \mid g^t\right)\Phi_{t+1} m_{t+1} + \Phi_t = 0 \tag{24}$$

$$m_{t+1} : -\Gamma_t\left[V_{t+1} + \theta\left(1 + \ln m_{t+1}\right)\right] - \Phi_{t+1} + \omega_{t+1} = 0 \tag{25}$$

The $t = 0$ first order conditions, which are functions of the initial levels of capital and debt, are detailed in Appendix A.

There are two points worth noting about the set of optimality conditions. First, the first order conditions, and consequently the robustly optimal allocation, do not depend upon the level of the probability distortion, M_t. This result stems from the assumption that the government takes as its objective function the consumers' subjective expected utility.[5] Because the expectations of the two agents are aligned, the government does not attempt to use its policy tools to re-align the consumers' subjective expectation with the approximating probability model. Rather, the government sets its taxes to induce the best path for the allocation and probability distortions, taking as given the current level of consumer beliefs.

Second, (22) indicates that the multiplier Γ_t is a martingale under the subjective expectation. That is, $\widetilde{E}_{t-1}\Gamma_t = \Gamma_{t-1}$. A similar property is found in Svec (2011). This martingale affects the persistence of the allocation. In the limit as $\theta \to \infty$, the multiplier becomes constant over time and across states.

3.1.1 Ramsey policies and prices:

The solution to the Ramsey problem yields the equilibrium allocation and probability distortions. The bond holdings in each state, then, are given by (19). Given these values, this section describes the policies and prices that implement the solution. That is, using the solutions that come from the Ramsey problem, the goal of this section is to determine the prices $\{w, r\}$, bond returns $\{R_b\}$, and taxes $\{\tau, \Omega\}$ that decentralize the equilibrium. To accomplish this goal, I use the consumer's budget constraint and the first order conditions from the consumer's and the firm's problems.

The prices on capital and labor follow directly from the competitive firm's marginal product conditions. The labor tax rate can then be determined through the consumer's intra-temporal condition:

$$\tau_t = 1 + \frac{u_l(c_t, l_t)}{u_c(c_t, l)\, F_l(k_{t-1}, l_t, g_t)} \tag{26}$$

Thus, the intra-temporal wedge is uniquely pinned down by the allocation.

[5] If, instead, the planner maximizes the expected utility of the consumers with respect to the approximating model, then the allocation would be a function of the distortion, M_t.

The two remaining variables to find are R_b and Ω. The equations used to determine these values at time $t+1$ are

$$1 = \beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) m_{t+1} \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)} R_{b,t+1}, \tag{27}$$

$$1 = \beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) m_{t+1} \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)} R_{k,t+1} \tag{28}$$

where

$$R_{k,t+1} = 1 + \left(1 - \Omega_{t+1}\right)\left(r_{t+1} - \delta\right), \tag{29}$$

and the $t+1$ consumer's budget constraint:

$$c_{t+1} + k_{t+1} + b_{t+1} - \left(1 - \tau_{t+1}\right) w_{t+1} l_{t+1} - R_{b,t+1} b_t - R_{k,t+1} k_t = 0. \tag{30}$$

As this set of equations makes clear, there are more unknowns than equations. Consequently, this model cannot separately identify R_b and Ω. To see this, suppose that there are N states of the world at time $t+1$. This means that there are $2N$ variables that must be pinned down and only $N+2$ equations. This indeterminacy is worsened by the fact that there is one additional linear dependency among the constraints. This can be seen by multiplying (30) by $\beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) m_{t+1} u_c\left(c_{t+1}, l_{t+1}\right)$ and by summing the result over g_{t+1}. The outcome is a function only of the allocation and distortions and not $R_{b,t+1}$ or $R_{k,t+1}$. Thus, model uncertainty does not overturn the indeterminacy of the capital tax rates and debt returns, as found by Chari, Christiano, and Kehoe (1994).

Because of this indeterminacy, the state-contingent capital tax rates and bond returns cannot be separately identified. However, the theory pins down two policy variables related to these instruments. The first instrument is the ex-ante capital tax rate, defined as

$$\Omega_t^e \equiv \frac{\sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) m_{t+1} \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)} \Omega\left(g^{t+1}\right)\left[F_{k,t+1} - \delta\right]}{\sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) m_{t+1} \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)}\left[F_{k,t+1} - \delta\right]}. \tag{31}$$

This ex-ante capital tax rate is the consumers' subjective expectation of the $t+1$ capital tax rate, weighted by the stochastic discount factor. Using (13), the numerator can be shown to equal

$$\sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) m_{t+1} \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)}\left[F_{k,t+1} + 1 - \delta\right] - \frac{1}{\beta}, \tag{32}$$

which is a function entirely of the allocation. Consequently, the ex-ante capital tax rate can be determined. This ex-ante value is different from the version in Chari, Christiano, and Kehoe (1994) in that the expectation is taken with respect to the subjective probability model, rather than with the approximating model.

The second policy variable pinned down by the theory is labeled the private assets tax rate because it combines information from both the ex-post capital tax rate and the return on government debt. To

derive this variable, suppose that the debt return in each state in period $t+1$ is the combination of a non-state-contingent return and a state-contingent tax rate:

$$R_{b,t+1} = 1 + \bar{r}_t \left[1 - \nu_{t+1}\right], \tag{33}$$

where the non-state-contingent rate of return, $\bar{r}_t$, must satisfy

$$\sum_{g_{t+1}} \pi\left(g_{t+1} \mid g_t\right) m_{t+1} \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)} R_{b,t+1} = \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g_t\right) m_{t+1} \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)} \left[1 + \bar{r}_t\right]. \tag{34}$$

This constraint implies that

$$\sum_{g_{t+1}} \pi\left(g_{t+1} \mid g_t\right) m_{t+1} \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)} \nu_{t+1} = 0. \tag{35}$$

With this decomposition, the non-state-contingent return on debt can be determined through (12). From the government's budget constraint, the total tax revenues from capital and debt in a particular state g_{t+1}

$$\Omega_{t+1}\left[r_{t+1} - \delta\right] k_t + \nu_{t+1} \bar{r}_t b_t \tag{36}$$

are equal to

$$g_{t+1} - \tau_{t+1} w_{t+1} l_{t+1} - b_{t+1} + \left(1 + \bar{r}_t\right) b_t. \tag{37}$$

Finally, in order to turn this value into a rate and ease comparisons to the ex-ante capital tax rate, divide by the total return across capital and bonds in each state. Then, the private assets tax rate is

$$\eta_{t+1} = \frac{\Omega_{t+1}\left[r_{t+1} - \delta\right] k_t + \nu_{t+1} \bar{r}_t b_t}{\left[r_{t+1} - \delta\right] k_t + \bar{r}_t b_t}. \tag{38}$$

Overall, this fiscal policy model with capital pins down the wage, the rental rate of capital, and three tax variables: a labor tax, the ex-ante capital tax, and a private assets tax. In order to determine the specific characteristics of these prices and policies, we construct the recursive version of the planner's optimization problem and solve it numerically. But, before we follow this procedure, there is one policy result that can be analytically derived by focusing attention on a specific, and simple, class of functions describing the consumers' preferences. We highlight this implication in the following section.

3.1.2 Ex-ante capital tax rate under preference restrictions:

A powerful finding of Chari, Christiano, and Kehoe (1994) is that, within a specific class of utility functions, the ex-ante capital tax rate is exactly equal to zero. However, one might fear that this policy conclusion hinges upon the assumption that consumers have rational expectations. In this section, we re-examine whether this theoretical implication still survives when consumers face model uncertainty.

For this section, assume that the utility function of the consumers is quasi-linear, where

$$u(c,l) = c + v(l) \tag{39}$$

Plugging this functional form into the consumer's first order condition with respect to capital for $t > 0$, the equation becomes

$$1 = \beta \sum_{g_{t+1}} \pi(g_{t+1} \mid g_t) m_{t+1} \{1 + (1 - \Omega_{t+1})(F_k(k_t, l_{t+1}, g_{t+1}) - \delta)\} \tag{40}$$

The planner's first order condition with respect to the same variable is

$$1 = \beta \sum_{g_{t+1}} \pi\left(g_{t+1} \mid g^t\right) m_{t+1} \frac{\mu_{t+1}}{\mu_t} [F_k(k_t, l_{t+1}, g_{t+1}) + 1 - \delta] \tag{41}$$

where μ_t is the Lagrange multiplier on the resource constraint in period t. Combining these two equations with (20), the numerator of the ex-ante capital tax rate is equal to

$$\beta \sum_{g_{t+1}} \pi(g_{t+1} \mid g_t) m_{t+1} \left[\frac{\Gamma_t - \Gamma_{t+1}}{\Gamma_t - 1 - \xi}\right] [F_k(k_t, l_{t+1}, g_{t+1}) + 1 - \delta] \tag{42}$$

Proposition 2 *$\forall t > 1$, if $cov_t\{(\Gamma_t - \Gamma_{t+1}), (F_{k,t+1} + 1 - \delta)\} = 0$ under the consumer's subjective expectation, then $\Omega_t^e = 0$. $\Omega_t^e \neq 0$ otherwise.*

Proof. *To see this, first note that Γ_t is a martingale under the consumer's subjective expectation, where $\underset{\sim}{E}_t \Gamma_{t+1} = \Gamma_t$. Then, a property of covariance suggests that the numerator is equal to*

$$\frac{\beta}{\Gamma_t - 1 - \xi} cov_t\{(\Gamma_t - \Gamma_{t+1}), (F_{k,t+1} + 1 - \delta)\}$$

It follows from (31) that $\Omega_t^e = 0$ only when $cov_t\{(\Gamma_t - \Gamma_{t+1}), (F_{k,t+1} + 1 - \delta)\} = 0$ and $\Omega_t^e \neq 0$ when this condition does not hold. ■

This proposition provides a simple test to determine whether the value of the ex-ante capital income tax rate is equal to 0 for a given value of θ. In the limit as $\theta \to \infty$, the Lagrange multiplier Γ_t is constant across time $\Gamma_t = \Gamma, \forall t$. This implies that the covariance is equal to zero and hence the ex-ante capital tax rate is also equal to 0. This is the case examined by Chari, Christiano, and Kehoe (1994). Outside of this limit, though, the covariance is no longer equal to zero, meaning that the ex-ante capital tax rate is also non-zero.

Intuitively, this result stems from the fact that the planner must consider how its choice of capital taxes affects the consumers' incentive to save as well as their endogenous beliefs. This second desire can be seen through the first term in the covariance: $\Gamma_t - \Gamma_{t+1}$. This random variable tracks the shadow value of the consumers' welfare across states, which, in turn, reflects the consumers' probability distortion across those same states. In balancing these two incentives, the government allows the shadow value of the consumers' welfare, Γ, to fluctuate.

Another perspective confirms the logic underlying the proposition. It can be shown that if

$$\frac{\mu_{t+1}}{\mu_t} = \frac{u_c\left(c_{t+1}, l_{t+1}\right)}{u_c\left(c_t, l_t\right)} \tag{43}$$

then $\Omega_t^e = 0, \forall t \geq 1$. That is, if the planner places the same value on resources over time as the consumer, then the ex-ante capital tax rate is equal to 0. This condition is satisfied in a rational expectations model. However, when consumers face model uncertainty, the planner values resources differently than the consumers. This is because the planner, when considering whether to allocate more consumption to the consumers in one state, takes into account not just the consumers' marginal utility gain from that action, but also the effect that action has on the consumers' probability distortion. It is this additional marginal value that breaks the equality in (43). Thus, there is no theoretical presumption that the ex-ante capital tax rate is equal to 0, even under quasi-linear preferences.[6]

3.2 Recursive formulation of the ramsey problem:

This section describes the recursive formulation of the planner's problem. Government spending is now assumed to follow a Markov process. The natural state vector is a function of both capital and government spending. However, because of the forward-looking constraint on the movement of the consumers' subjective welfare, V_t, this problem is not time-consistent. As detailed by Marcet and Marimon (1998), the addition of a co-state variable allows this constraint to be written recursively. The co-state variable, Γ, keeps track of the past promises made by the planner about the consumers' subjective welfare.

The time 0 values of the capital stock, debt, and probability distortion imply that the period 0 problem of the government is unlike the problem it faces in all other periods. To account for this difference, the recursive formulation has to be separated into two. The first Bellman equation presented below applies to the planner's problem in any period $t > 0$, while the second one applies only to $t = 0$. When calculating the path of the economy over time, the values of the endogenous variables coming from the $t = 0$ problem will be used as inputs into the $t > 0$ problem.

The planner's value function, $H\left(\cdot; \xi\right)$ satisfies the following Bellman equation:

[6] Although I have written the proof assuming a quasi-linear form of consumer preferences, a similar argument can be made for a utility function of the following form:

$$u\left(c, l\right) = \frac{c^{1-\sigma}}{1-\sigma} + v\left(l\right)$$

The only difference is that (42) would contain the ratio $\frac{u_{c,t+1}}{u_{c,t}}$ in the expectation, which, in turn, would modify the covariance term in the proof.

$$H\left(k_{-},\Gamma_{-},g_{-};\xi\right)=\min_{\Gamma_g}\max_{c_g,l_g,V_g,k_g,m_g}\sum_{g}\pi\left(g\mid g_{-}\right)\left\{\begin{array}{c} m_g u\left(c_g,l_g\right)+\xi m_g\left[u_c\left(c_g,l_g\right)c_g+u_l\left(c_g,l_g\right)l_g\right] \\ +m_g\mu_g\left[c_g+g+k_g-F\left(k_{-},l_g,g\right)-\left(1-\delta\right)k_{-}\right] \\ -\Gamma_{-}\left[m_g V_g+\theta m_g\ln m_g\right]+m_g\Gamma_g\left[V_g-u\left(c_g,l_g\right)\right] \\ +\omega_g\left[m_g-\frac{\exp\left(\frac{-V_g}{\theta}\right)}{\sum_{g}\pi(g|g_{-})\exp\left(\frac{-V_g}{\theta}\right)}\right]+\beta m_g H\left(k_g,\Gamma_g,g;\xi\right) \end{array}\right\}$$

There are many points worth noting here. First, this Bellman equation is written from an ex-ante perspective. This formulation is necessary because of the presence of the incremental probability distortion. As noted above, this distortion is a function of the characteristics across all states within the same time period. In order to capture this, the Bellman equation must be expressed before the realization of uncertainty. Thus, the subscript g denotes the state-contingent value of each random variable.

Second, the solution to recursive portion of the planner's problem is indexed by the multiplier ξ on the planner's implementability constraint. For a given ξ, the first order conditions and additional constraints imply an optimal allocation in periods one and beyond. The implied recursive policy functions for a given ξ must, therefore, be checked for consistency with time-zero allocations and the resulting implementability constraint. Our numerical algorithm solves simultaneously for recursive policies and the value of ξ that satisfies implementability given those policies.

The time 0 recursive problem of the planner is

$$H_0=\min_{\Gamma_0}\max_{c_0,l_0,V_0,k_0}\left\{\begin{array}{c} u\left(c_0,l_0\right)+\xi\left[u_c\left(c_0,l_0\right)c_0+u_l\left(c_0,l_0\right)l_0\right]-\xi u_c\left(c_0,l_0\right)\left[R_{b,0}b_{-1}+R_{k,0}k_{-1}\right] \\ +\mu_0\left[c_0+g_0+k_0-F\left(k_{-1},l_0,g_0\right)-\left(1-\delta\right)k_{-1}\right] \\ +\Gamma_0\left[V_0-u\left(c_0,l_0\right)\right]+\beta H\left(k_0,\Gamma_0,g_0\right) \end{array}\right\}$$

where $R_{k,0}=1+\left(1-\Omega_0\right)\left(F_k\left(k_{-1},l_0,g_0\right)-\delta\right)$. The first order conditions for both of these recursive problems are detailed in Appendix B. There, they are verified to be equivalent to those derived in the sequential formulation of the Ramsey problem.

4 Calibration and numerical solution method:

In the calibration we draw, as much as possible, on the parameters used by Chari, Christiano, Kehoe (1994). Preferences take the form

$$u(c_t,l_t)=\frac{\left(c_t^{1-\gamma}l_t^{\gamma}\right)^{\Psi}}{\Psi}\quad . \tag{44}$$

We assume throughout that $\beta=0.98$, $\gamma=0.75$. We take log-utility ($\Psi=0$) as our baseline, and consider the robustness of our results to parameterization with higher risk-aversion ($\Psi=-5$). The capital share is

given by $\alpha = 0.34$, and capital depreciation by $\delta = 0.08$. Government spending is assumed to follow the following autoregressive process:

$$\log(g_t/\bar{g}) = \rho_g \log(g_{t-1}/\bar{g}) + \epsilon_t. \tag{45}$$

We take $\bar{g} = 0.07$, which corresponds to a steady-state government expenditure share of about 16 percent of GDP. We consider a baseline i.i.d. case, $\rho_g = 0$, and a persistent case, $\rho_g = 0.89$. For each calibration of the autocorrelation parameter, we select σ_ϵ to match an unconditional variance of government expenditure of 15 percent. Initial period parameters are also fixed so as to be consistent with Chari, Christiano, Kehoe (1994). We fix $k_{-1} = 1.05$, $R_{b0}b_{-1} = 0.2$, and $\Omega_0 = 0.271$. Finally, we always assume the economy begins with initial government spending g_0 equal to $\bar{g}$, its long-run average.

To solve the model, we discretize the autoregressive process for government expenditure according to the method advocated by Kopecky and Suen (2010). We then solve the model using a collocation approach. The recursive policy functions are approximated by weighted sums of linear finite-element basis functions over a simplictical grid on the endogenous states, k and Γ. We then use a numerical equation solver to simultaneously solve for the basis-function weights that satisfy the recursive first order conditions at the grid points as well as the initial-period allocations and value of ξ that satisfy the corresponding period-zero first order conditions.

5 Findings:

This section describes and interprets the optimal policy chosen by the government when faced with consumer uncertainty. In calculating this optimal policy, we initially assume that the consumers have logarithmic preferences ($\Psi = 0$) and that the fiscal shock is not persistent ($\rho = 0$). We label this model our baseline model. After describing and interpreting the optimal policy in the baseline model, we run two robustness checks to determine how sensitive our results are to changing parameter values. The first robustness check increases the consumers' risk aversion, and the second robustness check increases the persistence of the fiscal shock.

Optimal policy involves one period of transition and then, for all future periods, stationary policies that are functions of an augmented state vector. During the period of transition, the labor tax is negative and the capital income tax is highly positive for all levels of consumer uncertainty. In Figure 1, we plot initial period taxes as a function of consumer uncertainty, assuming the baseline model's parameter values.[7] In that figure, the values of the initial taxes are plotted on the y-axis and the size of the distortion associated with consumer uncertainty is plotted on the x-axis. Small levels of distortion represent little consumer

[7]Throughout, we plot outcome as a function of the total distortion as measured by Φ_0 rather than the corresponding value of θ. This choice allows for comparability of results across different specification of preferences and the exogenous processes.

uncertainty, while large levels of distortion represent high consumer uncertainty. Also, as the distortion shrinks to zero, the economy returns to the benchmark, rational expectations equilibrium described in Chari, Christiano, and Kehoe (1994). Figure 1 shows that consumer uncertainty does not substantially change the transition-period taxes.

Tables 1-3 report our main results from our baseline model about how consumer uncertainty affects the stationary policies chosen by the government. These tables describe the time series properties of the government's policy tools for three different values of consumer uncertainty.[8]

In Table 1, we can see that when consumers face little uncertainty (corresponding to "no distortion"), the three main results of Chari, Christiano, and Kehoe (1994) hold: the ex-ante capital income tax rate is zero, the labor tax is large and smooth across states, and the private assets tax is highly volatile. As consumer uncertainty rises and we move to Tables 2 and 3, we see how increasing consumer uncertainty affects optimal policy. Specifically, greater consumer uncertainty seems to have three main policy implications. First, as uncertainty rises, the government sets an increasingly negative ex-ante capital tax rate. This numerical finding is consistent with our analytical result from above that the expected capital income tax is no longer identically equal to zero, even under logarithmic preferences. That being said, the government chooses a very modest capital subsidy. Under our baseline parameterization, the largest subsidy (which occurs when consumer uncertainty is greatest) is approximately 0.022%. Evidently, consumer uncertainty does not provide a significant justification for allowing the ex-ante capital income tax rate to diverge far from zero.

The second policy implication is that the government should increase the persistence of the labor tax with the size of the consumers' uncertainty. To see this, note that while the persistence of the spending shock remains constant, the autocorrelation of the labor tax rises from 0.54 when there is no distortion to 0.99 when there is high distortion. This increase in persistence raises the labor tax's standard deviation and reduces the correlation of the labor tax with government spending. Thus, high levels of consumer uncertainty seem to impart a random walk component to the labor tax. This finding is consistent with our analytical result that consumer uncertainty makes the labor tax a function of the Lagrange multiplier on the consumers' value function, Γ, which is itself a martingale. The last policy implication of consumer uncertainty is that the government increases both the mean and the volatility of the private assets tax as uncertainty rises.

In Figure 2, we plot the time series properties of the fiscal variables against the size of the distortion associated with consumer uncertainty. As that figure makes clear, the qualitative policy implications discussed above remain true: greater uncertainty leads the government to (modestly) subsidize capital,

[8]While the table only displays three different values of consumer uncertainty, we have solved the numerical model for many more values of θ and have found that our qualitative results are consistent with Table 1.

make labor taxes more persistent, and increase both the mean and the volatility of the private assets tax.

To examine our results in more depth, we now turn to the model's impulse response functions. These impulse responses compare the government's chosen policies across two different time paths of the exogenous shock. The first shock process assumes that government spending equals its average level for all periods except one, at which point government spending rises to $g = g_h$. The second shock process assumes that government spending always equals its average level. Using this information, we then calculate the percentage change in each variable across the two shock processes at each point in time. The results are displayed in Figure 3.

Figure 3 shows that, in response to a rise in government spending, consumer uncertainty doesn't substantially affect the ex-ante capital tax rate, but it does lead to qualitative differences in the optimal labor tax. Compared to the rational expectations benchmark, the labor tax under consumer uncertainty rises by less and remains persistently lower. Qualitatively, this means that not only does consumer uncertainty lead to a smaller jump in the labor tax on the shock's impact, but that the government permanently adjusts the labor tax in response to a rise in public expenditures. Given this profile of labor taxes, the government must turn to the private assets tax to finance both the immediate increase in spending as well as the persistently lower labor tax. As a consequence, the private assets tax rises by even more than is optimal under the benchmark model. With these additional funds, the government purchases assets from the consumers, which it will then use to finance the future reduction in labor tax revenue. These additional assets show up in the impulse response function for bonds, as we can see that the value of bonds is negative under high consumer uncertainty. Effectively, the government finances the permanent reduction in labor taxes through a one-time increase in the private assets tax.

Taken together, we can start to understand how consumer uncertainty affects the optimal policy response to fluctuations in government spending. Upon impact of a rise (fall) in spending, consumer uncertainty leads the government to keep both the labor tax and the ex-ante capital income tax largely unchanged, and instead raise the sensitivity of the private assets tax. This larger change in the private assets tax allows the government to purchase more assets (issue more debt), which will then be used to finance the permanently lower (higher) value of the labor tax.

These results have implications for how much of the fiscal shock is absorbed by the labor and private assets taxes, where shock absorption for a particular policy instrument is defined as the increase in tax revenue obtained by that instrument divided by the rise in government spending cumulated over the period of the impulse response. In Figure 4, we depict the shock absorption of both the labor income tax and the private assets tax. Consistent with our impulse response functions, we see that as consumer uncertainty grows, the labor tax absorbs less of the shock, while the private assets tax absorbs more. In fact, for high levels of uncertainty, the private assets tax absorbs more than 100% of the entire government

spending shock, which then allows the government to purchase the additional assets which help pay for the permanent reduction in the labor tax.

At this point, we understand *how* the government sets fiscal policy when facing consumer uncertainty, but the question remains as to *why* the government would choose these particular policies. In the following few paragraphs, we argue that there are two main motivations behind the planner's choices, both of which exploit the endogeneity of the consumers' expectations. We discuss each motivation in turn.

The first motivation of the planner is to manipulate the price of the state-contingent public debt in order to reduce the cost of its fiscal insurance. Before we describe how the government achieves this cost reduction, though, it would be helpful to better understand what is meant by fiscal insurance. It is well known in the optimal fiscal policy literature with state-contingent public debt that it is optimal for the planner to purchase insurance from the consumers against the public spending shock; Lucas and Stokey (1983) and Chari, Christiano, and Kehoe (1994) are two examples of this result. To do this, the government purchases assets that pay off when government spending is high and issues debt that must be repaid when government spending is low. Under the insurance interpretation, the consumers effectively receive their insurance premium when government spending is low, but must pay out damages (accept lower returns on their private assets) when government spending is high. The benefit of this insurance is that it allows the government to maintain a relatively smooth profile of labor taxes across states.

Now, to this benchmark model, the current paper inserts consumer uncertainty. The consumers, in their uncertainty about the shock process, endogenously tilt their subjective probability model away from the approximating model. As can be seen in (10), the resulting subjective probability model depends on the consumers' state-contingent welfare, values that depend on the policy choices of the government. Specifically, if the government uses its policy tools to increase (decrease) the consumers' welfare in state A, (10) suggests that all else equal the consumers would decrease (increase) their subjective likelihood placed on state A. This movement, in turn, would affect the stochastic discount factor, $\beta m_{t+1} \frac{u_{c,t+1}}{u_{c,t}}$, and so the price of public debt. In particular, if the consumers believe that state A is less (more) likely to occur, then one would expect that the price of the state A-contingent asset would fall (rise), raising (reducing) its return. As a consequence, in setting its fiscal policy, the government must consider how its policy choices influence the price of its fiscal insurance.

As consumer uncertainty now exposes an additional channel through which the government can affect the economy's asset prices, the government exploits this endogeneity to move prices in a beneficial direction, decreasing the price of the assets it wants to purchase (making its insurance cheaper) and increasing the price of the debt it wants to issue (making the debt it issues more dear). To accomplish the first change and lower the price of assets that pay off when government spending is high, the logic above suggests that the government must lower the labor tax in the high government spending state. This would increase

the consumers' welfare, reducing their subjective likelihood that this state will occur, and so decrease the asset's price. To accomplish the second change and raise the price of the debt it issues that is to be repaid when spending is low, the logic above suggests that the government must raise the labor tax in the low spending state. This would decrease the consumers' welfare, increasing their subjective likelihood that this state will occur, and so increase the asset's price. Thus, the first motivation of the government – the motivation to manipulate the economy's asset prices – leads to a smoother profile of labor taxes across states. A similar argument is made in Karantounias (2013) in a model without capital.

The second motivation is more straight-forward: the government wants to use its fiscal policy to smooth the consumers' welfare across states in order to mitigate the direct welfare costs of uncertainty aversion, a motivation that depends on the fact that the planner maximizes the consumers' *subjective* expected utility. To see this direct welfare cost, consider a policy choice that results in large fluctuations in the consumers' welfare across states. Using (10), we can see that the large fluctuations would lead to large probability distortions by the consumers. This means that the consumers' subjective probability model would be distant from the approximating model and that the consumers would place a much smaller (larger) weight on the high (low) welfare state. This probability tilting directly reduces the consumers' subjective expected utility. If, though, the government could change its policy to smooth the consumers' welfare across states, then (10) suggests that the probability distortions would be smaller. As a consequence, the consumers' uncertainty aversion would not lead to such a large fall in the consumers' subjective welfare.

Given this second motivation, the government wants to use its policy tools to smooth the consumers' welfare across states.[9] This means that the government wants to raise the consumers' welfare in the high spending state and lower the consumers' welfare in the low spending state. It accomplishes these movements by lowering the labor tax in the high spending state and raising the labor tax in the low spending state, the same policy implications as found in the first motivation. Thus, both motivations of the Ramsey planner are consistent: smooth the consumers' labor tax rate across states! These policy movements can be seen numerically in Figures 3 and 4.

The two motivations discussed above – the price-manipulation motivation and the welfare-smoothing motivation – also influence the desire of the planner to make its labor tax persistent. To see this, note that both motivations rely on the government manipulating the consumers' probability distortions, m: for the price-manipulation motivation, influencing m leads to beneficial asset prices, while for the welfare-smoothing motivation, the goal is to reduce the size of the distortions themselves, which mitigates the direct welfare cost of uncertainty aversion. These probability distortions, in turn, are functions of the consumers'

[9] As discussed earlier, an alternative assumption we considered was that the Ramsey planner could maximize the consumers' expected utility under the approximating probability model. Although we do not numerically solve this model in this paper, we predict that in this alternative model, the planner will still choose to smooth taxes across states because of motivation 1, but less aggressively because motivation 2 no longer holds.

state-contingent welfare, V, which depend on both the contemporaneous utility of the consumers and the discounted sum of their future utilities. This means that the government can influence the consumers' probability distortions by adjusting today's policies or by adjusting future policies. Given this, it seems intuitively clear that the government will use both margins to influence the consumers' behavior so that it doesn't need to make labor taxes highly sensitive to government spending (and so volatile). This logic implies that when the government has the incentive to reduce (raise) labor taxes in one period relative to the benchmark model, the incentive should persist into the future, leading to permanently lower (higher) labor taxes. We see exactly this type of persistence in the impulse response function in Figure 3.

To determine the robustness of our policy results, we now turn to sensitivity tests in which we sequentially vary two key parameter values: the consumers' risk aversion and the persistence of the government spending shock. We describe and interpret the resulting policies below.

The first robustness check we run involves increasing the consumers' risk aversion to $\Psi = -5$. In doing this, we modify θ appropriately to maintain the same overall distortion as measured by the discounted sum of relative entropy. Just as before, optimal policy involves a period of transition and then, for all future periods, policy becomes a stationary function of the augmented state vector. During that period of transition, the government again sets a negative labor income tax and a highly positive capital income tax. As in Chari, Christiano, and Kehoe (1994), greater risk aversion does substantially increase the initial ex-ante capital income tax for all levels of consumer uncertainty (it is approximately 786% in the baseline model and it is 1159% in the high risk aversion case). This movement makes sense because the larger is the consumers' risk aversion, the less distortionary is the capital income tax. This leads the government to rely more heavily on that initial capital tax to finance its future spending.

Tables 4-6 report the main properties of the stationary policies chosen by the government in the case of high risk aversion. We can see in Table 4 that high risk aversion does not change the three main consequences of consumer uncertainty for policy: consumer uncertainty still leads to a negative ex-ante capital income tax rate, greater persistence in the labor income tax, and a larger (on average) and more volatile private assets tax.

That being said, risk aversion does lead to quantitative differences in the optimal policy, if not qualitative differences. These differences stem from the fact that greater risk aversion implies that capital taxation is less distortionary. As mentioned above, greater risk aversion leads the government to implement a higher ex-ante capital income tax in the transition period. This greater stock of initial assets allows the government to increase its capital subsidy (on average) and to reduce the mean value of the labor income tax. Also, because the intertemporal substitutability of consumption is lower, there is less of a welfare cost associated with varying the capital tax. Consequently, the government allows the ex-ante capital income tax rate to fluctuate more than in the baseline model.

Figure 5 plots the impulse response functions of the economy in response to a one period rise in government spending. As we saw in the baseline model, these impulse response functions suggest that consumer uncertainty leads the government to reduce the labor income tax in the same period as the spending shock (relative to the rational expectations result) and keep this lower labor tax permanently. In fact, with greater risk aversion, the government lets labor taxes fall by more in response to the spending shock than is optimal in the baseline model. This is consistent with our findings from Tables 4-6. The government pays for this by increasing both the ex-ante capital income tax and the private assets tax by more than we saw in Figure 3. This greater influx of tax revenue means that the government purchases more assets from the consumers, allowing the government to pay for the permanently lower labor income tax. Finally, the movements in the labor and private assets tax imply that greater risk aversion leads the government to rely more heavily on the private assets tax and less heavily on the labor tax to finance shocks to public expenditures.

For our second robustness check, we increase the persistence of the government spending shock to $\rho = 0.89$. This change does not substantially impact the optimal policies chosen during the transition period, when the labor tax is -36% and the capital income tax is 786%. But, as we will see, the stationary optimal policies are qualitatively different than under the baseline model. We report the properties of the stationary policies in Tables 7-9, while the impulse response functions are plotted in Figure 6.

As indicated in Table 7, the increased persistence leads the government to increase the autocorrelation of the labor tax, lower the mean of the private assets tax, and raise its standard deviation. All of these changes are consistent with the results in Chari, Christiano, and Kehoe (1994).

Tables 8 and 9, meanwhile, indicate that the optimal policies chosen under consumer uncertainty are sensitive to the persistence of the fiscal shock. This makes sense because of the inherent pessimism associated with uncertainty aversion. That is, when consumers are uncertainty averse, they fear that the true probability model is one that hurts their welfare or, in this case, puts a high likelihood on positive fiscal shocks. Now, when government spending is persistent, the feared model is more damaging to the consumers in that they fear that the fiscal shock is persistently positive. As a consequence, the consumers' behavior should be even more distorted than under an iid shock, which should lead the planner to respond more aggressively.

Looking at Tables 8 and 9, we see that greater consumer uncertainty leads the planner to lower the mean value of the private assets tax substantially. Now, for high levels of consumer uncertainty, the average private assets tax is approximately -11%, while it is only -0.3% in the baseline model. The government also chooses to dramatically raise the standard deviation of the private assets tax as consumer uncertainty increases. At the same time that the government reduces the private assets tax rate, it raises the mean value of the ex-ante capital tax. This movement runs counter to what we observe in the baseline

model, where greater uncertainty leads to a greater (though modest) capital subsidy. In fact, in the high distortion case, the ex-ante capital income tax is approximately 3%. The other noteworthy point about the ex-ante capital tax rate is that the government no longer holds the tax fixed as consumer uncertainty rises. Rather, the standard deviation of the tax rate is 13.5 percentage points at high levels of consumer uncertainty. Evidently, consumer uncertainty leads the planner to respond to a positive (negative) fiscal shock by reducing (raising) the ex-ante capital tax.

Turning to the impulse response functions in Figure 6 and focusing on the high consumer uncertainty case, we see that increased persistence markedly affects how the government responds to a fiscal shock. Now, in response to a positive fiscal shock, the government chooses to reduce the consumers' capital income tax. This response would likely spur the consumers to increase their savings. To finance this, the government contemporaneously raises both the labor tax and the private assets tax. Although the spending shock is more persistent in this economy, the impulse response function assumes that public expenditures fall back to their average levels after the one period shock. This surprise – that spending did not remain high for many periods – leads the government to reduce the labor tax and the private assets tax in period 2. That being said, the labor tax remains persistently higher because of the past positive shock to government spending. The movement in the labor tax is opposite of what we encountered in the baseline model, where high consumer uncertainty leads the government to permanently reduce the tax rate.

6 Conclusion:

This paper examines how consumer uncertainty affects the optimal policies implemented by a fiscal authority in a model with capital. Unlike in a rational expectations framework, consumers in this model are uncertain as to the true probability model governing the shock process. Wanting to be robust against this uncertainty, they apply a max-min operator to their decision problems. That is, the consumers choose the allocation that maximizes their expected utility, where the expectation is taken with respect to their subjective probability model. While the consumers face model uncertainty, we assume that the government is fully confident that the approximating model correctly characterizes the stochastic environment and yet for political economy reasons chooses to maximize the consumers' subjective expected utility.

Given these preferences, the government seeks to use its fiscal policy to mitigate the welfare costs associated with both the assumed linearity in the tax rates and consumer uncertainty. It is shown analytically that, under one condition, the government no longer implements a zero ex-ante capital tax rate, as is optimal within the rational expectations benchmark model. This is because the government takes into account how the consumers' allocation affects their probability distortion.

After numerically solving this model, we find that consumer uncertainty affects the optimal stationary policies chosen by the government in three ways. First, greater uncertainty leads the government to subsidize the consumers' capital income on average, although the size of the subsidy is quite modest. Second, uncertainty leads the government to increase the persistence of the labor tax, imparting a random walk component to the tax. Third, greater uncertainty leads the government to raise both the mean and the standard deviation of the private assets tax. This tax can be shown to absorb more of the government spending shock as uncertainty rises.

These changes are rationalized as the product of two motivations that are novel to the consumer uncertainty environment: a price-manipulation motivation and a welfare-smoothing motivation. The first motivation leads the government to choose policies that reduce the price of assets it wants to purchase and raise the price of assets that it wants to sell; intuitively, these actions reduce the cost to the government of purchasing fiscal insurance against its spending shock. The second motivation leads the government to choose policies that smooth the consumers' welfare across states because this alleviates the direct welfare costs associated with uncertainty aversion.

7 References:

1. Aiyagari, R. (1995). Optimal capital income taxation with incomplete markets, borrowing constraints, and constant discounting. *Journal of Political Economy* 103(6), 1158-1175.

2. Bossaerts, P., P. Ghirardato, S. Guarnaschelli, and W. Zame (2010). Ambiguity in Asset Markets: Theory and Experiment. *The Review of Financial Studies* 23(4), 1325-1359.

3. Chamley, C. (1986). Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives. *Econometrica* 54 (3), 607-622.

4. Chari, V. V., L. Christiano, and P. Kehoe (1994). Optimal fiscal policy in a business cycle model. *Journal of Political Economy* 102 (4), 617-652.

5. Conesa, J., S. Kitao, and D. Krueger (2009). Taxing capital? Not a bad idea after all! *American Economic Review* 99(1), 25-48.

6. Correia, I. (1996). Should capital income be taxed in the steady state? *Journal of Public Economics* 60, 147-151.

7. Dennis, R. (2010). How robustness can lower the cost of discretion. *Journal of Monetary Economics 57*(6), 653-667.

8. Dennis, R., K. Leitemo, and U. Soderstrom (2009). Methods for robust control. *Journal of Economic Dynamics and Control 33*(8), 1604-1616.

9. Guo, J. and K. Lansing (1999). Optimal taxation of capital income with imperfectly competitive product markets. *Journal of Economic Dynamics and Control* 23, 967-995.

10. Hansen, L. P. and T. J. Sargent (2001). Robust control and model uncertainty. *American Economic Review* 91: 60-66.

11. Hansen, L. P. and T. J. Sargent (2005). Robust estimation and control under commitment. *Journal of Economic Theory* 124 (2): 258-301.

12. Hansen, L. P. and T. J. Sargent (2007). Recursive robust estimation and control without commitment. *Journal of Economic Theory* 136, 1-27.

13. Hansen, L. P. and T. J. Sargent (2008). Robustness. Princeton University Press.

14. Imrohoroglu, S. (1998). A quantitative analysis of capital income taxation. *International Economic Review* 39(2), 307-328.

15. Judd, K. (1985). Redistributive Taxation in a Simple Perfect Foresight Model. *Journal of Public Economics* 71, 1-26.

16. Karantounias, A. (2013). Managing expectations and fiscal policy. *Theoretical Economics* 8, 193-231.

17. Klein, P. and J. Rios-Rull (2003). Time-consistent optimal fiscal policy. *International Economic Review* 44(4), 1217-1245.

18. Kopecky, K. and R. Suen (2010). Finite State Markov-chain Approximations to Highly Persistent Processes. *Review of Economic Dynamics,* 13(3), 701-714.

19. Leitemo, K. and U. Söderström (2008). Robust monetary policy in the new keynesian framework. *Macroeconomic Dynamics* 12, 126-135.

20. Levin, A. and J. Williams (2003). Robust monetary policy with competing reference models. *Journal of Monetary Economics* 50, 945-975.

21. Lucas, R. and N. Stokey (1983). Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics* 12, 55-93.

22. Marcet, A. and R. Marimon (1998). Recursive contracts. Universitat Pompeu Fabra Working paper 337.

23. Onatski, A. and J. Stock (2002). Robust monetary policy under model uncertainty in a small model of the U.S. economy. *Macroeconomic Dynamics* 6, 85–110.

24. Piketty, T. and E. Saez (2012). *A Theory of Optimal Capital Taxation.* NBER working paper 17989.

25. Stockman, D. (2001). Balanced-budget rules: welfare loss and optimal policies. *Review of Economic Dynamics* 4, 438-459.

26. Svec, J. (2011). Optimal fiscal policy with robust control. Forthcoming in *Journal of Economic Dynamics and Control.*

27. Walsh, C. (2004). Robustly optimal instrument rules and robust control: an equivalence result. *Journal of Money, Credit, and Banking* 36(6), 1105-1113.

28. Woodford, M. (2010). Robustly optimal monetary policy with near-rational expectations. *American Economic Review* 100(1), 274–303.

8 Appendix A:

The $t = 0$ first order conditions for a government that maximizes the consumers' expected utility under the distorted probability model are

$$\begin{aligned} c_0 \quad : \quad & 0 = u_c(c_0, l_0) + \xi \left[u_{cc}(c_0, l_0) c_0 + u_c(c_0, l_0) + u_{cl}(c_0, l_0) l_t \right]_0 + \mu_0 - \Gamma_0 u_c(c_0, l_0) \\ & -\xi u_{cc}(c_0, l_0) \left[R_{b0} b_{-1} + R_{k0} k_{-1} \right] \end{aligned}$$

$$\begin{aligned} l_0 \quad : \quad & 0 = u_l(c_0, l_0) + \xi \left[u_{cl}(c_0, l_0) c_0 + u_{ll}(c_0, l_0) l_0 + u_l(c_0, l_0) \right] - \mu_0 F_l(k_{-1}, l_0, g_0) - \Gamma_0 u_l(c_0, l_0) \\ & -\xi u_{cl}(c_0, l_0) \left[R_{b0} b_{-1} + R_{k0} k_{-1} \right] - \xi u_c(c_0, l_0)(1 - \Omega_0) F_{lk}(k_{-1}, l_0, g_0) \end{aligned}$$

$$V_t : 0 = \Gamma_0$$

$$k_0 : 0 = \mu_0 - \sum_{g_1} \beta \pi(g_1 \mid g_0) m_1 \mu_1 \left[F_k(k_0, l_1, g_1) + 1 - \delta \right]$$

9 Appendix B:

The first order conditions from the recursive formulation of the planner's problem are

$$c_g : 0 = u_c(c_g, l_g) + \xi\left[u_{cc}(c_g, l_g)\, c_g + u_c(c_g, l_g) + u_{cl}(c_g, l_g)\, l_g\right] + \mu_g - \Gamma_g u_c(c_g, l_g)$$

$$l_g : 0 = u_l(c_g, l_g) + \xi\left[u_{cl}(c_g, l_g)\, c_g + u_l(c_g, l_g) + u_{ll}(c_g, l_g)\, l_g\right] - \mu_g F_l(k_-, l_g, g) - \Gamma_g u_l(c_g, l_g)$$

$$k_g : 0 = \mu_g + \beta H_k(k_g, \Gamma_g, g; \xi)$$

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

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

$$\Gamma_g : 0 = V_g - u(c_g, l_g) + \beta H_\Gamma(k_g, \Gamma_g, g; \xi)$$

The envelope conditions are

$$H_k(k_-, \Gamma_-, g_-; \xi) = -\sum_g \pi(g \mid g_-)\, \mu_g m_g\left[F_k(k_-, l_g, g) + 1 - \delta\right]$$

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

Table 1: Ramsey tax moments for the basline calibration with iid shocks and no distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	22.661	0.032	0.539	0.803
Ex-ante Capital Tax	-0.000	0.000	0.889	0.061
Priv. Asset Tax	-0.733	39.268	0.033	0.998
Cap. Tax w. Safe Debt	-0.501	31.272	0.033	0.998

Table 2: Ramsey tax moments for the basline calibration with iid shocks and modest distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	22.634	0.045	0.756	0.580
Ex-ante Capital Tax	-0.005	0.000	0.894	0.073
Priv. Asset Tax	-0.420	41.248	0.034	0.998
Cap. Tax w. Safe Debt	-0.247	32.883	0.034	0.998

Table 3: Ramsey tax moments for the basline calibration with iid shocks and high distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	22.546	0.163	0.986	0.128
Ex-ante Capital Tax	-0.022	0.001	0.911	0.106
Priv. Asset Tax	0.714	48.776	0.033	0.997
Cap. Tax w. Safe Debt	0.676	39.049	0.033	0.997

Table 4: Ramsey tax moments for the risk-averse calibration with i.i.d. shocks and no distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	19.802	0.001	0.465	0.821
Ex-ante Capital Tax	-0.180	2.578	0.019	0.998
Priv. Asset Tax	1.443	42.931	0.166	0.982
Cap. Tax w. Safe Debt	1.007	25.304	0.167	0.983

Table 5: Ramsey tax moments for the risk-averse calibration with i.i.d. shocks and modest distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	19.764	0.147	0.986	-0.073
Ex-ante Capital Tax	-0.185	2.560	0.020	0.998
Priv. Asset Tax	2.456	55.516	0.142	0.986
Cap. Tax w. Safe Debt	1.636	33.026	0.144	0.986

Table 6: Ramsey tax moments for the risk-averse calibration with i.i.d. shocks and high distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	19.833	0.694	0.986	-0.079
Ex-ante Capital Tax	-0.258	2.575	0.023	0.995
Priv. Asset Tax	4.539	102.558	0.092	0.990
Cap. Tax w. Safe Debt	3.104	62.095	0.094	0.991

Table 7: Ramsey tax moments for the basline calibration with persistent shocks and no distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	22.827	0.114	0.888	0.999
Ex-ante Capital Tax	0.000	0.000	0.879	-0.998
Priv. Asset Tax	-10.123	141.397	-0.007	0.450
Cap. Tax w. Safe Debt	-6.960	109.608	-0.003	0.457

Table 8: Ramsey tax moments for the basline calibration with persistent shocks and modest distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	22.846	0.187	0.935	0.742
Ex-ante Capital Tax	0.150	1.376	0.879	-0.998
Priv. Asset Tax	-10.818	152.882	-0.013	0.437
Cap. Tax w. Safe Debt	-7.283	118.200	-0.009	0.446

Table 9: Ramsey tax moments for the basline calibration with persistent shocks and high distortion.

	Mean	**Std. Dev.**	**Autocorr.**	**Corr. w/ G**
Labor Tax	23.737	0.994	0.965	0.523
Ex-ante Capital Tax	2.997	13.525	0.877	-0.998
Priv. Asset Tax	-11.062	218.613	-0.043	0.347
Cap. Tax w. Safe Debt	-5.980	184.219	-0.037	0.366

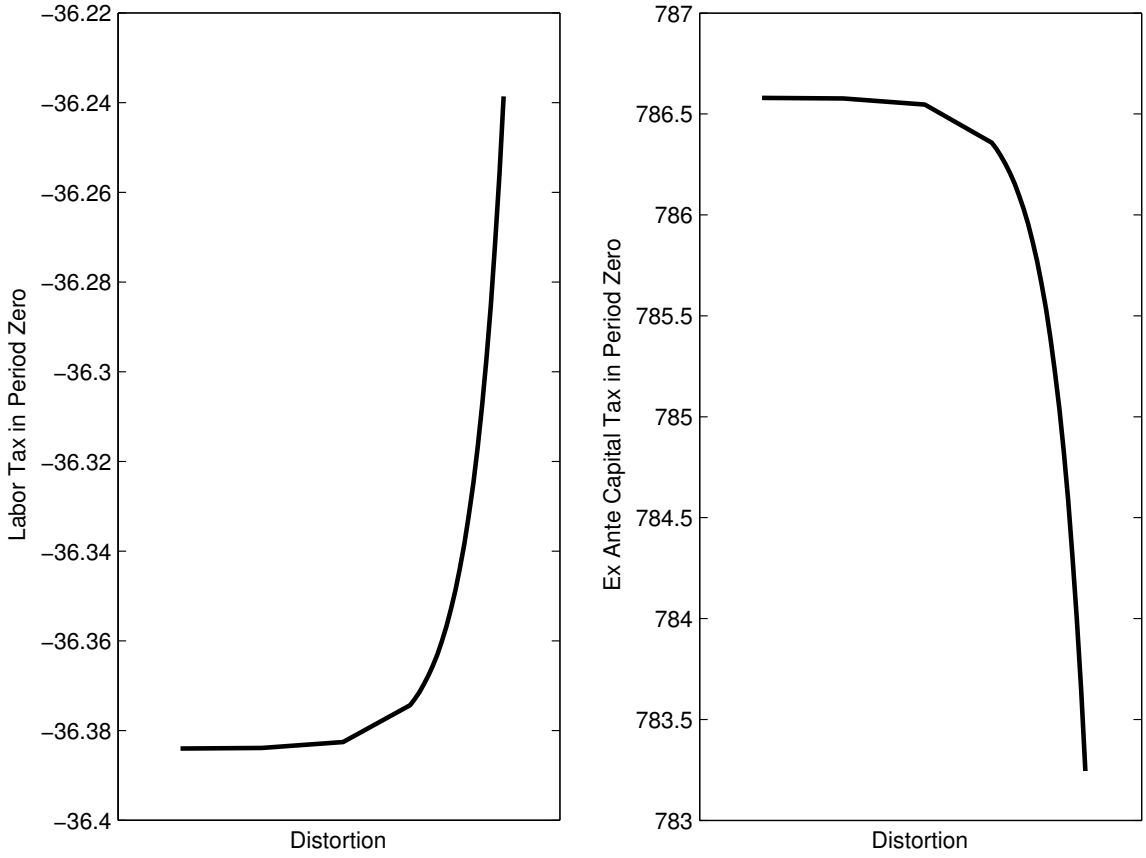

Figure 1: Period zero tax rates under baseline parameterization and i.i.d. government spending shocks.

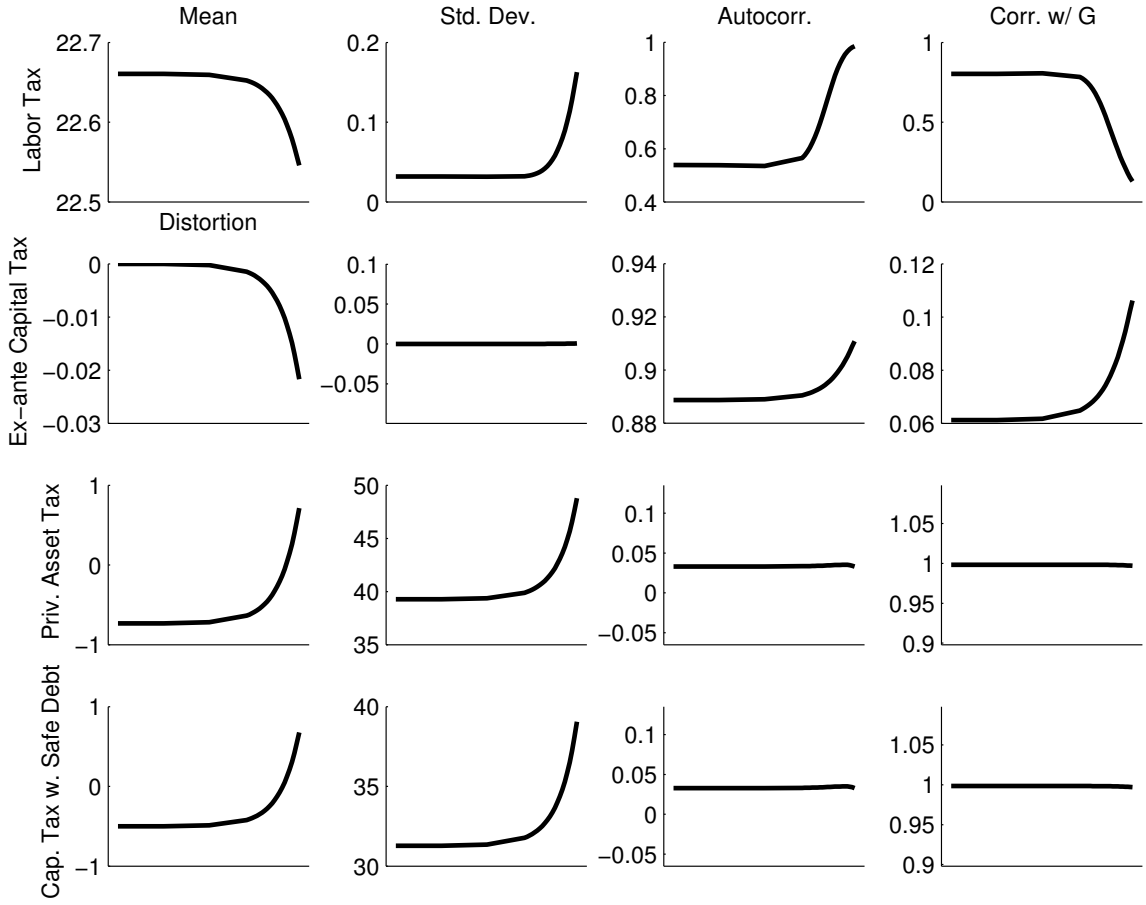

Figure 2: Fiscal instrument moments as a function of the degree of distortion in consumers' beliefs, for the baseline parameterization with i.i.d. government spending shock.

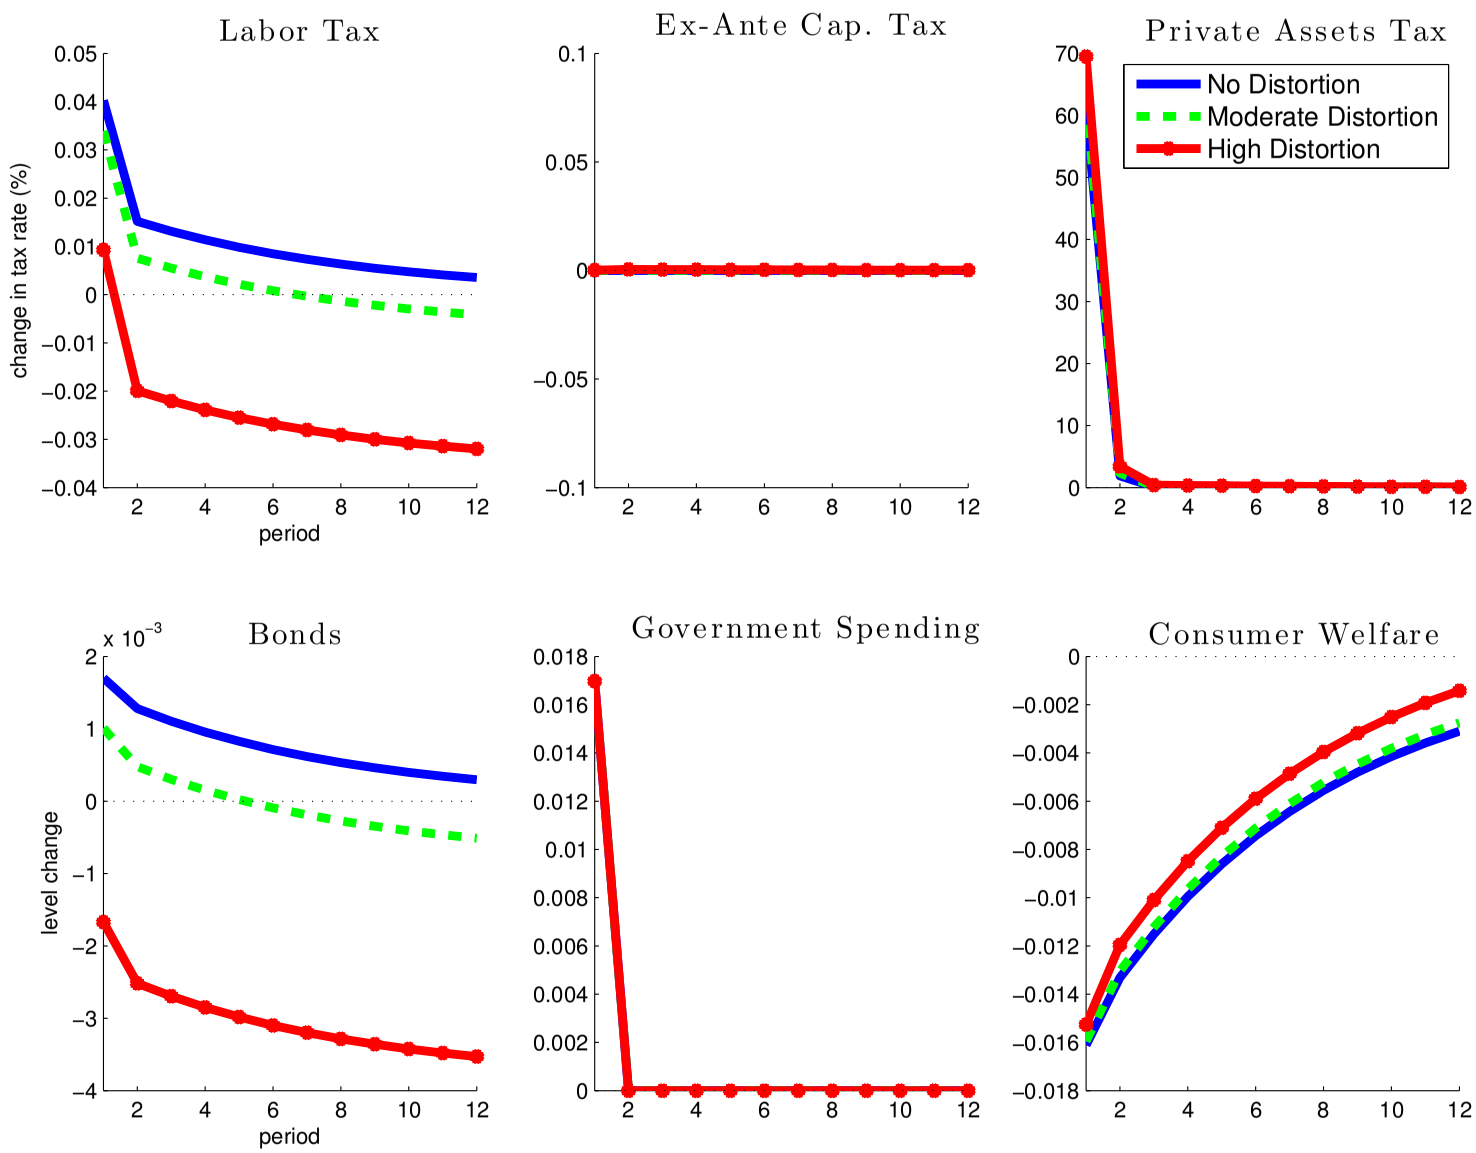

Figure 3: Impulse responses to a one-time deviation of government expenditure, for the baseline parameterization with i.i.d. government spending shocks.

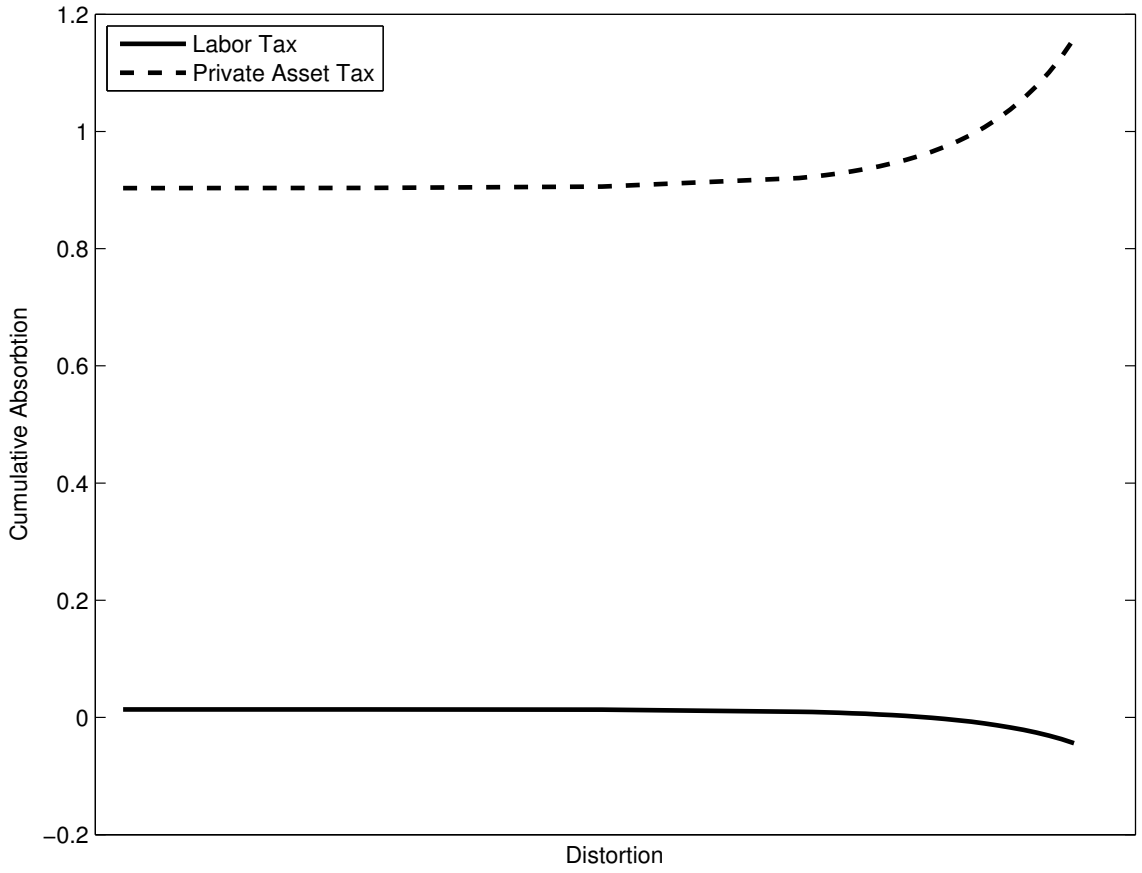

Figure 4: Labor and asset tax absorption as a function of the distortion in consumers' beliefs, for the baseline parameterization with i.i.d. government spending shocks.

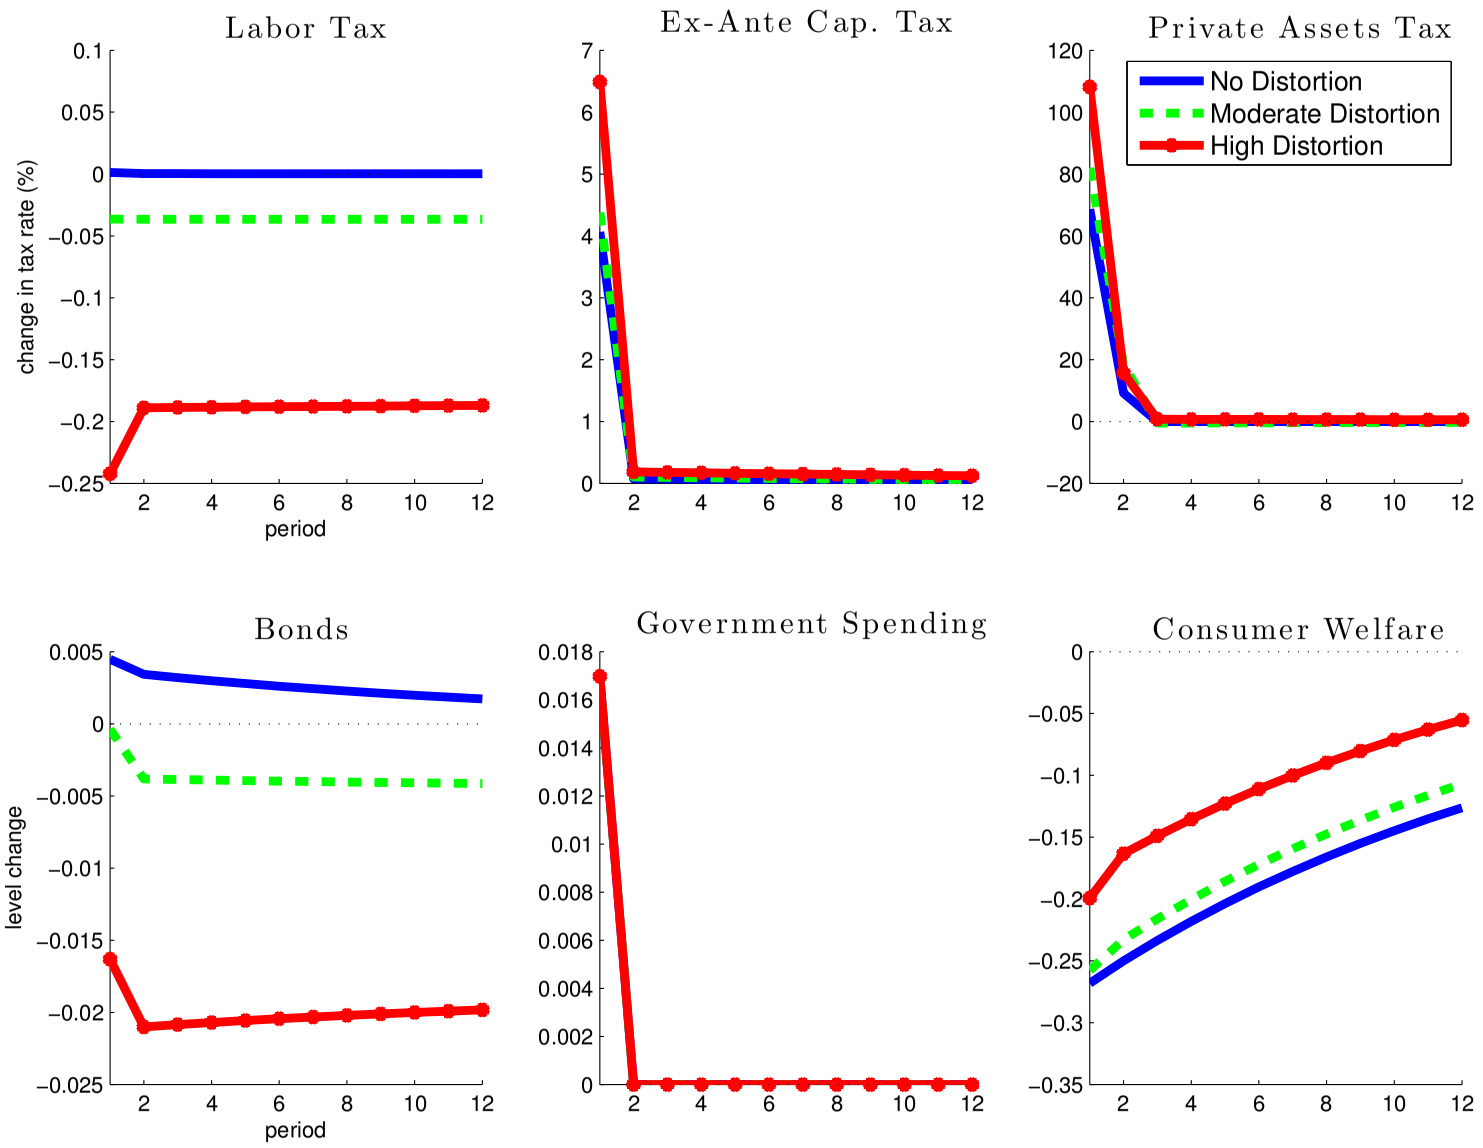

Figure 5: Impulse responses to a one-time deviation of government expenditure, for the high risk aversion ($\Psi = 5$) parameterization with i.i.d. government spending shocks.

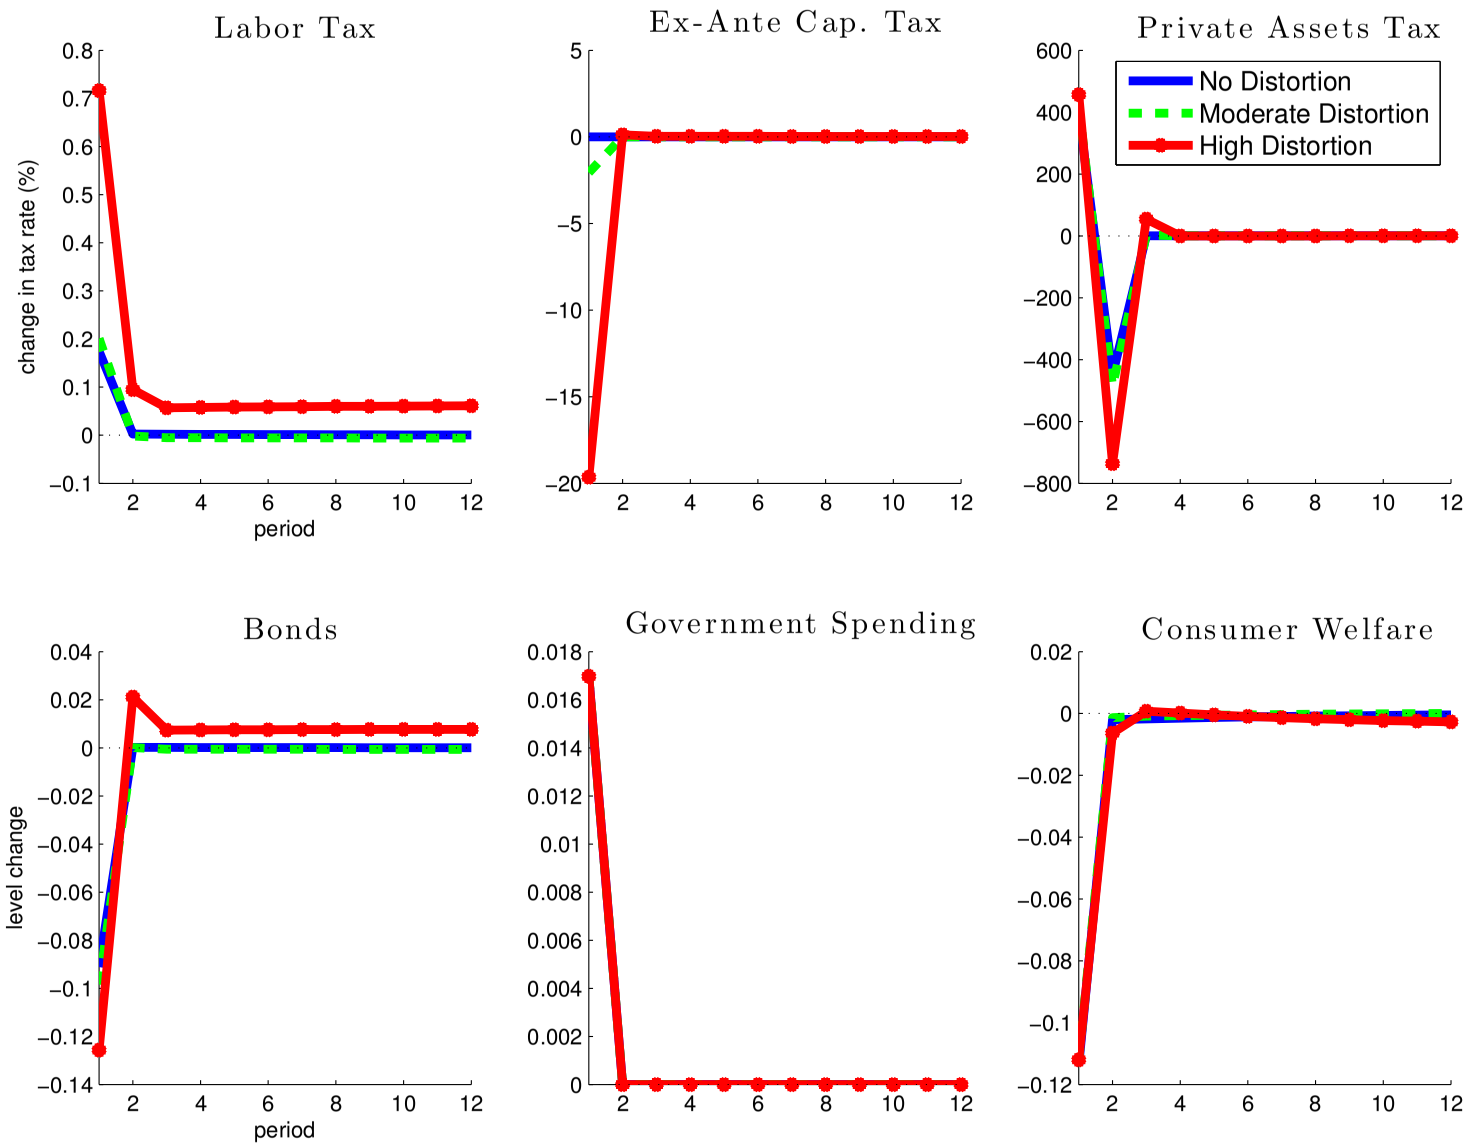

Figure 6: Impulse responses to a one-time deviation of government expenditure, for the baseline parameterization (log utility) with persistent government spending shocks.