Auburn University
Department of Economics
Working Paper Series

The Consequences of Uncertain Debt Targets

Alexander Richter[†], and Nathaniel Throckmorton[*]

[†]*Auburn University*, [*]*Indiana University*

AUWP 2013-18

The Consequences of Uncertain Debt Targets*

Alexander W. Richter Nathaniel A. Throckmorton

October 24, 2013

ABSTRACT

Recent proposals to reduce U.S. debt reveal large differences in their implied targets. These differences demonstrate the uncertainty surrounding future tax rates and long-run debt targets. We use a standard real business cycle model in which a Bayesian household learns about the state-dependent debt target in an endogenous tax rule. The household extracts the debt target state from a noisy tax process and jointly estimates the transition probabilities. We compare the household's ability to learn and the consequences of the uncertainty across different limited information sets. Limited information influences the household's behavior but also imposes two-sided risk. Despite the popular viewpoint that fiscal uncertainty has negative effects, limited information can result in welfare gains or losses, depending on whether the household's expectations are consistent with the realization of future states. Although the welfare distribution includes gains, we stress that the uncertainty created by the recent fiscal policy debate slowed the current recovery and led to welfare losses. If Congress provides clarity about future policy, output and welfare increase and the economy quickly recovers.

Keywords: Bayesian learning; Limited information; Fiscal uncertainty; Welfare
JEL Classifications: D83; E32; E62; H68

*Richter, Department of Economics, Auburn University, 0332 Haley Center, Auburn, AL (arichter@auburn.edu); Throckmorton, Department of Economics, Indiana University, 100 S. Woodlawn, Wylie Hall 105, Bloomington, IN (nathrock@indiana.edu). We especially want to thank Eric Leeper and Todd Walker for their advice and feedback. We also thank Randy Beard, Toni Braun, Tim Conley, Bill Gavin, Grey Gordon, Bulent Guler, Juan Carlos Hatchondo, Amanda Michaud, Gilad Sorek, conference participants at the 2013 Midwest Economic Association Meetings, the Spring 2013 Midwest Macroeconomics Meetings, and the 2013 Computational Economics and Finance Meetings, and seminar participants at Indiana University, the St. Louis Federal Reserve, and the Atlanta Federal Reserve.

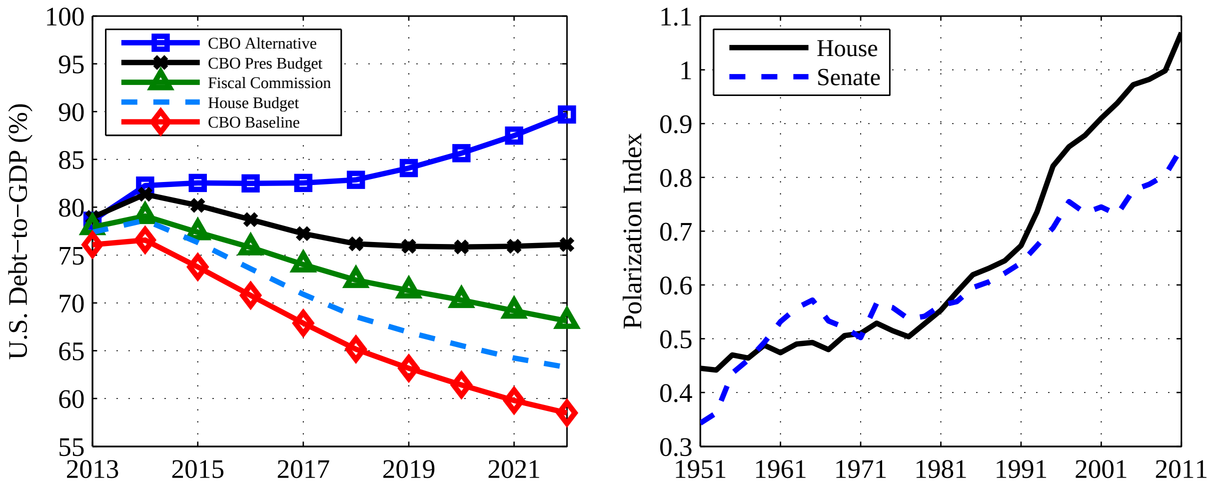

(a) Long-run debt targets implied by some of the U.S. debt reduction proposals, compared with the baseline (current law) and alternative fiscal financing scenarios in the CBO's 2012 Long-Term Budget Outlook. Source: Committee for a Responsible Federal Budget (2011)

(b) Index of polarization of the U.S. Congress, measuring the difference between the two parties' positions on a conservative-liberal interval, $[-1, 1]$. The minimum of this index is 0 (full agreement) and the maximum is 2 (full polarization). Source: Poole and Rosenthal (2012)

Figure 1: Measures of fiscal uncertainty

1 INTRODUCTION

Recent proposals to reduce U.S. debt reveal large differences in their implied 10-year targets, with estimates ranging from 63 to 76 percent of GDP in 2022 (figure 1a). The CBO's projections show an even wider variance. The differences in these estimates and their inconsistency with the post-WWII average debt-to-GDP percentage, which is roughly 45 percent, demonstrate the potential for sharp swings in policy and the uncertainty surrounding future tax rates and long-run debt targets.

What makes the fiscal outlook dangerous and the degree of fiscal uncertainty unprecedented? Political polarization has risen steadily since the 1950s and is currently at a record level (figure 1b). High polarization reduces the likelihood of compromise and weakens the signal households receive from Congress. The deficit reduction proposals provide evidence that Congress is aware of the structural deficit and its potential consequences, but without a credible plan that commits to a debt target, households must form expectations over a range of debt targets, which affects their beliefs about future tax rates, their optimal consumption/saving decisions, and welfare.[1]

This paper uses a novel application of Bayesian learning to endogenize fiscal uncertainty within a standard real business cycle model. Each year the fiscal authority sets a state-dependent debt target and a corresponding income tax rate consistent with the target, the debt-to-output ratio, and discretionary tax policy. The household knows the tax rule and observes tax rates and debt-to-GDP ratios, but does not know the current debt target because it is obscured by discretionary tax policy. The household rationally learns about the debt target with the Hamilton (1989) filter, and uses

[1]During WWII many industrialized countries faced unprecedented debt-to-GDP percentages. For example, the debt-to-GDP percentage in the U.S., the U.K., Australia, Canada, and New Zealand was 122, 270, 92, 115, and 148 percent. Over the next 50 years (and in many cases sooner), these percentages all fell below 60 percent and in most cases well-below 50 percent [Abbas et al. (2011)]. There are two reasons for the sharp reduction. First, all of these countries ran primary surpluses on average after the war. Second, and more importantly, average inflation exceeded 3.5 percent, which reduced real interest rates, and growth rates ranged from 2.5-5 percent. History suggests large debt reductions are possible, but given the current polarization and anemic growth, debt reduction is very uncertain.

observations of the tax rate and debt-to-output ratio to adaptively learn the debt target transition matrix with a Gibbs sampler [Albert and Chib (1993)]. Given the household's sequence of past beliefs about the debt target, each period they estimate the transition matrix most likely to generate the sequence and optimally solve for their decision rules conditional on their estimate. When the household's beliefs are incorrect, they are surprised by future tax rates that are inconsistent with their expectations. These surprises distort the real economy, since the tax rate is levied on income.

To quantify the consequences of uncertain debt targets, we simulate the model with different information sets. The benchmark is the full information set, where current and past debt target states and the transition matrix is known. In this case, there is no learning and the only sources of uncertainty are future states and discretionary changes in the tax rate. We compare three limited information sets to this benchmark: 1. the debt target state is always unknown and the transition matrix is known; 2. current and past states are known and the transition matrix is unknown; and 3. both the state and transition matrix are always unknown. In case 1, the household incorporates their state likelihood estimates into their expectations and rationally learns about the state as time evolves. In case 2, the household uses Bayesian methods to estimate the transition matrix, but they treat their estimate as the truth when forming expectations. This has two implications. First, the household adaptively learns about the transition matrix (i.e., they are boundedly rational). Second, it introduces Knightian uncertainty; the household guesses the risk based on current beliefs, but is unable to accurately quantify it without the true distribution. Case 3 is a hybrid of cases 1 and 2.

Conditional on each limited information set, we find there are persistent deviations from the full information (case 0) paths. When the household faces their most limited information set (case 3), the paths of capital and output regularly deviate from the case 0 paths by 0.3 and 0.1 percent. Despite the popular viewpoint that fiscal uncertainty has negative economic effects,[2] these deviations do *not* necessarily lead to welfare losses, since the household faces two-sided risk. The key determinant of welfare is whether future tax rates are more or less extreme than what the household expects. If the household places greater weight on a higher debt target, meaning lower expected taxes, then they invest more than they would under full information. If the true debt target winds up being even higher (lower), the household gains (loses). Thus, there is an equal chance of welfare gains and losses when the household's prior and the true transition matrix are symmetric.

Although a portion of the welfare distribution shows gains, these outcomes did not recently occur since the realizations of policy were inconsistent with expectations. If we focus on the period from 2010-2012, people were uncertain if the Bush tax cuts would be extended and placed positive probability on both outcomes. In reality, the Bush tax cuts were extended and taxes wound up lower than people expected. This means their investment strategy was suboptimal. Had they known the tax cuts would be extended, they would have invested more, given the higher after-tax returns. In our model, we conduct an experiment that shows this policy debate slowed the recovery and reduced welfare. If Congress provides greater clarity about the future direction of fiscal policy, even after a period of uncertainty, output and welfare increase and the economy quickly recovers.

The remainder of the paper is organized as follows. Section 2 connects our work with the literature. Section 3 introduces the economic model, the information sets, and the numerical methods. Section 4 describes the calibration and connects our model to the data. Section 5 presents

[2]See, for example, *The Wall Street Journal*, "Want Growth? Try Stable Tax Policy" (December 21, 2011), *The Economist*, "The cloud of uncertainty. Dithering in the dark." (June 16, 2012) and "Pointless, painful uncertainty" (July 18, 2012), and *The New York Times*, "Uncertainty Is the Enemy of Recovery" (April 28, 2013).

the results. We first illustrate the basic channels of the uncertainty and how incorrect inferences affect the household's decisions. Then we show how quickly the household learns and the welfare consequences of the uncertainty using model simulations. Section 6 concludes.

2 CONTACTS WITH THE LITERATURE

Several of the early papers that study policy uncertainty in a dynamic general equilibrium model permit discrete shifts in policy, so that agents place positive probability on recurring regime changes [Aizenman and Marion (1993); Bizer and Judd (1989); Dotsey (1990)]. A key feature in some of the more recent papers is that the current regime is unobserved [Andolfatto and Gomme (2003); Bianchi and Melosi (2013); Davig (2004); Leeper and Zha (2003); Schorfheide (2005)]. Instead, agents rationally form expectations about the regime using Bayesian updating. We build on this literature by also assuming agents do not know the transition probabilities governing regime change.

Much of the least squares learning literature models uncertainty in a similar way [see Evans and Honkapohja (2001, 2009) for an overview]. Contrary to rational expectations, agents do not have complete information about the structure of the economy and policy behavior, and instead act as econometricians using historical data to form beliefs about unknown components. For example, Giannitsarou (2006) finds that when households have limited information about the structure of the economy, a capital tax cut is ineffective during a recession. In our approach, the household knows the structure of the economy, but is uncertain about the state and transition matrix. They filter a noisy signal using the correct model and rationally form expectations, knowing the state is unknown. However, they use an anticipated utility approach [Kreps (1998)] to form expectations conditional on their transition matrix belief. This means the household adaptively forms expectations, conditional on a history of observations and their limited information set.

A segment of the rational expectations literature models fiscal uncertainty by assuming that households form expectations over a wide range of possible resolutions to the looming fiscal crisis [Bi et al. (2013); Davig and Foerster (2013); Davig and Leeper (2011); Davig et al. (2010, 2011); Richter (2012)]. This work makes clear that expectations about future policy, and the possible resolutions agents condition on, significantly affect equilibrium outcomes. The main drawbacks with this literature is that it requires the modeler to take a strong stance on the potential resolutions and it assumes the household knows the current policy state and the true probability distributions.

Another segment of the rational expectations literature uses stochastic volatility shocks to model fiscal uncertainty. Fernandez-Villaverde et al. (2011) and Born and Pfeifer (2013) both find that the impact of fiscal uncertainty is small. In their baseline calibrations, a two standard deviation volatility shock to various fiscal instruments decreases output by 0.06 and 0.025 percent. Johannsen (2013) finds that fiscal uncertainty shocks have a relatively large economic effect when the nominal interest rate is pegged at zero, but a much smaller effect when it is positive. These results are consistent with our findings, despite the fact that in our model the uncertainty is about the debt target state and its evolution and not time-varying variances. We find that a two standard deviation positive tax shock reduces output by 0.03 percent when the debt target state is hidden (case 1). However, when the transition probabilities are also unknown (case 3), output differs by 0.15 percent from its full information path. Moreover, we caution that these results do not exclude scenarios where fiscal uncertainty has more significant negative implications.

Of course, tax uncertainty, and more generally, other types of fiscal uncertainty, are only one component of the more broadly defined term "economic uncertainty". Many have argued that uncertainty about the recovery has reduced demand and also prolonged the recession. Much of the stochastic volatility literature focuses on precisely this question by examining how uncertainty about productivity and the business cycle affects economic conditions. The literature is mixed on the impact of these shocks. Some studies find that a one standard deviation volatility shock reduces output [Alexopoulos and Cohen (2009); Basu and Bundick (2012); Bloom (2009); Bloom et al. (2012)], while others report little or no impact [Bachmann and Bayer (2013); Chugh (2012)]. The stochastic volatility approach to modeling uncertainty provides an excellent first pass at quantifying the consequences of high and low periods of uncertainty. The drawback is that it provides little insight about the precise cause of the uncertainty. In our model, the source of uncertainty stems from limited information about the debt target, which is grounded in the recent policy debate.

3 Economic Model, Expectations, and Numerical Methods

We adopt a real business cycle model to study the consequences of uncertain debt targets. Our innovation is to introduce an expectations operator conditional on the household's beliefs about the current state and its transition matrix. The benchmark case, defined by the full information set, is when the household knows current and past states and the true transition matrix. We compare this case with the household's behavior conditional on three different limited information sets.

3.1 Real Business Cycle Model A representative household chooses $\{c_j, n_j, k_j, i_j, b_j\}_{j=t}^{\infty}$ to maximize expected lifetime utility, given by,

$$\mathbb{E}_t^{\ell} \sum_{j=t}^{\infty} \beta^{j-t} \left\{ \frac{c_j^{1-\sigma}}{1-\sigma} - \chi \frac{n_j^{1+\eta}}{1+\eta} \right\}, \tag{1}$$

where $\beta \in (0,1)$ is the discount factor, $1/\sigma$ is the intertemporal elasticity of substitution, $1/\eta$ is the Frisch elasticity of labor supply, c is consumption, n is labor hours, and $\mathbb{E}_t^{\ell}$ is an expectation operator conditional on information set ℓ (defined in section 3.2). The choices are constrained by

$$c_t + i_t + b_t = (1-\tau_t)(w_t n_t + r_t^k k_{t-1}) + r_{t-1} b_{t-1} + \bar{z}, \tag{2}$$

$$k_t = i_t + (1-\delta) k_{t-1}, \tag{3}$$

where w is the real wage rate, r^k is the real rental rate of capital, k is the capital stock, which depreciates at rate δ, b is a one-period real government bond with real return r, $\bar{z}$ is a fixed transfer from the government, and τ is a proportional tax rate levied against capital and labor income. The representative household's optimality conditions imply

$$w_t(1-\tau_t) = \chi n_t^{\eta} c_t^{\sigma}, \tag{4}$$

$$1 = \beta \mathbb{E}_t^{\ell} \{ (c_t/c_{t+1})^{\sigma} [(1-\tau_{t+1}) r_{t+1}^k + 1 - \delta] \}, \tag{5}$$

$$1 = \beta \mathbb{E}_t^{\ell} \{ (c_t/c_{t+1})^{\sigma} r_t \}. \tag{6}$$

A perfectly competitive representative firm produces output according to $y_t = a k_{t-1}^{\alpha} n_t^{1-\alpha}$, where a is technology and $\alpha \in (0,1)$. Each period the firm chooses $\{k_{t-1}, n_t\}$ to maximize profits, given by $y_t - w_t n_t - r_t^k k_{t-1}$. The optimality conditions imply

$$r_t^k = \alpha y_t / k_{t-1}, \tag{7}$$

$$w_t = (1-\alpha) y_t / n_t. \tag{8}$$

The fiscal authority finances constant discretionary spending, $\bar{g}$, and government transfers, $\bar{z}$, by levying taxes on income and issuing one-period real bonds. The government's flow budget constraint is given by

$$b_t + \tau_t(w_t n_t + r_t^k k_{t-1}) = r_{t-1} b_{t-1} + \bar{g} + \bar{z}. \tag{9}$$

The tax rate endogenously responds to fluctuations in the lagged debt-to-output ratio according to

$$\tau_t = \bar{\tau}(s_t) + \gamma(b_{t-1}/y_{t-1} - \overline{by}(s_t)) + \varepsilon_t, \tag{10}$$

where $\bar{\tau}(s_t)$ is a state dependent intercept of the tax rule, consistent with the stationary equilibrium implied by the current long-run debt target, $\overline{by}(s_t)$.[3] The tax shock, $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, is a proxy for discretionary tax policy (e.g., tax cuts during recessions). The debt target evolves according to an m-state Markov process with transition matrix, P. For row i and column j of P, element $p_{ij} = \Pr\{s_t = j | s_{t-1} = i\}$, for $i, j \in \{1, \dots, m\}$, where $0 \leq p_{ij} < 1$ and $\sum_{j=1}^m p_{ij} = 1$ for all i.

The aggregate resource constraint is given by

$$c_t + i_t + \bar{g} = y_t. \tag{11}$$

The transversality conditions omit explosive sequences for capital and debt and are given by

$$\lim_{T\to\infty} \beta^T \mathbb{E}_t\{(c_T)^{-\sigma} k_T\} = 0 \quad \text{and} \quad \lim_{T\to\infty} \beta^T \mathbb{E}_t\{(c_T)^{-\sigma} b_T\} = 0.$$

A competitive equilibrium consists of prices, $\{w_j, r_j^k, r_j\}_{j=t}^\infty$, quantities, $\{c_j, n_j, k_j, i_j, y_j\}_{j=t}^\infty$, government policies, $\{\tau_j, b_j, \bar{g}, \bar{z}\}_{j=t}^\infty$, and exogenous sequences, $\{\varepsilon_j, s_j\}_{j=t}^\infty$, which satisfy the household's and firm's optimality conditions, the government's budget constraint, the fiscal policy rules, the asset, labor, and goods markets' clearing conditions, and the transversality conditions.

3.2 EXPECTATIONS FORMATION

The different information sets are summarized in table 1. A formal description of the alternative information sets begins by writing the model compactly as

$$\mathbb{E}\left[f(\mathbf{v}_{t+1}^\ell, \mathbf{v}_t^\ell) | \Omega_t^\ell\right] = 0,$$

where f is a function that represents the structure of the model that is known to the household,

$$\mathbf{v}_t^\ell \equiv \begin{cases} (k_{t-1}, b_{t-1}, r_{t-1}, y_{t-1}, c_t, n_t, i_t, \tau_t, w_t, r_t^k, s_t), & \text{for } \ell \in \{0, 2\}, \\ (k_{t-1}, b_{t-1}, r_{t-1}, y_{t-1}, c_t, n_t, i_t, \tau_t, w_t, r_t^k, \mathbf{q}_t), & \text{for } \ell \in \{1, 3\}, \end{cases}$$

and Ω_t^ℓ is the household's information set in case $\ell \in \{0, 1, 2, 3\}$. In cases 1 and 3, $\mathbf{q}_t$ is a vector of conditional probabilities that $s_t = j$, for $j \in \{1, \dots, m\}$, which the household must update to form expectations. With $m = 3$, the (required) information sets are defined as follows,

$$\Omega_t^0 \equiv \{\Theta, \mathbf{z}_{t-1}^0, P\} \quad \text{(Case 0)} \qquad \Omega_t^2 \equiv \{\Theta, \mathbf{z}_{t-1}^2, \hat{P}_t, \mathbf{s}^t\} \quad \text{(Case 2)}$$
$$\Omega_t^1 \equiv \{\Theta, \mathbf{z}_{t-1}^1, P, \mathbf{x}^t\} \quad \text{(Case 1)} \qquad \Omega_t^3 \equiv \{\Theta, \mathbf{z}_{t-1}^3, \hat{P}_t, \mathbf{x}^t\} \quad \text{(Case 3)}$$

[3] While there is no consensus on the specification of the fiscal policy rule, a tax rule that endogenously responds to government debt is common. It is supported, in part, by the empirical findings of Bohn (1998), who contends that the U.S. primary surplus is an increasing function of the U.S. debt-to-GDP ratio. Davig and Leeper (2006), and references therein, estimate alternative fiscal rules and provide further evidence.

	Full Information		Limited Information	
	Case 0	Case 1	Case 2	Case 3
Debt Target State	Known	Unknown	Known	Unknown
Debt Target Transition Matrix	Known	Known	Unknown	Unknown

Table 1: Alternative information sets

where $\Theta \equiv \{\sigma, \eta, \chi, \delta, a, \alpha, \gamma, \bar{\tau}(1), \bar{\tau}(2), \bar{\tau}(3), \overline{by}(1), \overline{by}(2), \overline{by}(3), \sigma_\varepsilon^2\}$ is the set of model parameters shared by all information sets, $\mathbf{x}^t \equiv \{x_i\}_{i=0}^t$ and $\mathbf{s}^t \equiv \{s_i\}_{i=0}^t$ are histories of observations required for the Hamilton filter, importance sampler, and Gibbs sampler, $\hat{P}_t$ is the time-t estimate of the transition matrix from either the importance (case 2) or Gibbs (case 3) sampler and $\mathbf{z}_{t-1}^\ell \equiv \{k_{t-1}^\ell, r_{t-1}^\ell b_{t-1}^\ell, \tau_t^\ell\}$ is the minimum set of variables that initialize the problem in case ℓ.

Expectations are formed in one of two ways. In cases 0 and 2, the household knows the debt target state. Given $s_t = i$ and $\ell \in \{0, 2\}$, their expectations formation is

$$\mathbb{E}\left[f(\mathbf{v}_{t+1}^\ell, \mathbf{v}_t^\ell)|\Omega_t^\ell\right] = \sum_{j=1}^{3} p_{ij}^\ell \int_{-\infty}^{+\infty} f(\mathbf{v}_{t+1}^\ell, \mathbf{v}_t^\ell)\phi(\varepsilon_{t+1})\mathrm{d}\varepsilon_{t+1} = 0, \tag{12}$$

where p_{ij}^ℓ is the probability of transitioning to $s_{t+1} = j$, given information set ℓ and $\phi(\cdot)$ is the normal probability density function. Note that $\tau_{t+1} \in \mathbf{v}_{t+1}^\ell$ is a function of ε_{t+1}.

In cases 1 and 3, the discretionary tax shocks obscure the long-run debt target state and impose a signal extraction problem that the household solves to form beliefs about the state. The household observes the current tax rate and lagged debt-to-output ratio, and with their prior belief about the transition matrix applies the Hamilton (1989) filter to

$$x_t \equiv \tau_t - \gamma b_{t-1}/y_{t-1} = \bar{\tau}(s_t) - \gamma\overline{by}(s_t) + \varepsilon_t,$$

which is (10) solved for the combined observation, x_t. In other words, the household Bayesian updates the conditional probabilities, $\Pr(s_t = j|\mathbf{x}^t)$, given the previous conditional probabilities, $\Pr(s_{t-1} = i|\mathbf{x}^{t-1})$, with the most recent observation, x_t. Given these updated probabilities, $\mathbf{q}_t(i) \equiv \Pr(s_t = i|\mathbf{x}^t)$, and $\ell \in \{1, 3\}$, their expectation formation is

$$\mathbb{E}\left[f(\mathbf{v}_{t+1}^\ell, \mathbf{v}_t^\ell)|\Omega_t^\ell\right] = \sum_{i=1}^{3} \mathbf{q}_t(i) \sum_{j=1}^{3} p_{ij}^\ell \int_{-\infty}^{+\infty} f(\mathbf{v}_{t+1}^\ell, \mathbf{v}_t^\ell)\phi(\varepsilon_{t+1})\mathrm{d}\varepsilon_{t+1} = 0. \tag{13}$$

Including $\mathbf{q}_t$ in the state smooths expectations since it is unlikely that $\mathbf{q}_t(i) = 1$ for any i. Intuitively, the household knows that they do not know the state and is less responsive to new observations of x_t (i.e., they are cautious relative to case 0).

3.3 ALGORITHM AND SIMULATION PROCEDURE The model contains discrete changes in the debt target, which cause large deviations from the initial steady state. Thus, we solve the nonlinear model with a fixed-point projection method (linear interpolation and Gauss-Hermite quadrature). This solution technique simultaneously solves for the optimal policy functions at each point in the discretized state space. More specifically, it assumes the current policy functions hold at time t and $t + 1$ and uses the Euler equations to back out the updated policy functions. In this case, the number of states, $m > 2$, so the restriction $\sum_{i=1}^{m} \mathbf{q}_t(i) = 1$ imposes a constraint on the discretized state space. See appendix A.1 for details on how we discretize the state space.[4]

In all cases, we initialize the simulation procedure by drawing a random sequence of T true debt target states, $\{s_t\}_{t=1}^{T}$, and discretionary tax shocks, $\{\varepsilon_t\}_{t=1}^{T}$. Each period the fiscal authority sets the tax rate with (10), given the debt-to-output ratio, the debt target state, and the discretionary tax shock. The current tax rate is always observed by the household. Simulating the model in case 0 is straightforward, given the solution where expectations are evaluated according to (12).

In case 1, the household's prior belief about the transition matrix equals the truth and their beliefs about the state are built into their expectations, given in (13). Simulating case 1 requires updating the probability distribution of the state, $\mathbf{q}_t$, conditional on $\mathbf{x}^t$ each period (see appendix A.2 for details). Thus, the household Bayesian learns about the debt target state as time evolves.

In case 2, the household estimates the transition matrix each period, but, in expectation, the transition matrix belief is time invariant. In this sense, we follow the anticipated utility approach of Kreps (1998).[5] The household believes their estimate of the transition matrix is correct when forming expectations and solving for their decision rules. They use importance sampling from a Dirichlet distribution conditional on current and past observations of the state, $\mathbf{s}^t$, to update their prior belief about the transition matrix, $\hat{P}$, (see appendix A.3 for details). Each period in the simulation the household re-optimizes, (i.e., solves for new policy functions) given their new belief of the transition matrix. At the beginning of the simulation, the prior transition matrix is uninformative, so the distribution of their initial state inferences is uniform across all simulations.

In case 3, the household uses Gibbs sampling to make sequential draws of the vector of debt target states and the transition matrix (obtained with the importance sampler) conditional on the observations, $\mathbf{x}^t$. Their estimate of the transition matrix, $\hat{P}$, is the average of the last half of the resulting chain (see appendix A.4 for details). As with case 2, the prior transition matrix is uninformative. Each period the household re-optimizes after updating their estimates of the probabilities. The household is Bayesian learning about the debt target state as well as using an anticipated utility approach to form expectations with respect to the transition matrix.

In all cases, the fiscal authority uses the updated debt-to-output ratio, which is based on the household's decisions, the current debt target state, and the discretionary tax shock to set the tax rate in the next period of the simulation. The household uses this new information when filtering and sampling to improve their understanding of the debt target state and how it evolves.

4 CALIBRATION AND DATA CONNECTIONS

We parameterize the model at an annual frequency to study learning and the consequences of uncertainty over several decades. The baseline calibration, given in table 2, is consistent with the RBC literature. The discount factor, β, is set to equal 0.9615, which corresponds to a 4% annual real interest rate. The annual depreciation rate, δ, is set to equal 10 percent and the cost share of capital, α, is set to equal 0.33. We set the coefficient of relative risk aversion, σ, to equal 1, implying log utility in consumption. The Frisch elasticity of labor supply, $1/\eta$, is set to equal 0.5, which is consistent with the value used by the CBO and the findings of Chetty et al. (2011).

[4] See Judd (1998) and Richter et al. (2013) for further details and examples of fixed-point projection methods.

[5] Putting the entire transition matrix distribution in the state is computationally infeasible. However, Cogley and Sargent (2008) show that the anticipated utility approach provides a good approximation for Bayesian decision making.

Discount Factor	β	0.9615	Leisure Preference Parameter	χ	20.389
Capital Depreciation Rate	δ	0.10	Fiscal Policy Rule Coefficient	γ	0.30
Cost Share of Capital	α	0.33	Discretionary Tax Shock Std. Dev.	σ_ε	0.02
Constant of Relative Risk Aversion	σ	1	Low Debt Target State	$\overline{by}(1)$	0.60
Frisch Elasticity of Labor Supply	$1/\eta$	0.5	Mid Debt Target State	$\overline{by}(2)$	0.75
Technology Level	a	1.705	High Debt Target State	$\overline{by}(3)$	0.90
Steady State Government Spending Share	$\bar{g}/\bar{y}$	0.08	Probability of Remaining in State	p_{ii}	0.80
Steady State Government Transfers Share	$\bar{z}/\bar{y}$	0.09	Probability of Switching States	$p_{ij}, i \neq j$	0.10

Table 2: Baseline calibration

The leisure preference parameter, χ, implies a steady state share of time spent working equal to 0.33 in the middle debt target state. The level of technology, a, is set so that output is equal to 1 in the middle debt target state. We fix χ and a in the alternative states and solve for the implied steady state values of taxes, labor, and output that are consistent with the long-run debt targets.

The ratios of government expenditures/output and transfers/output are set to equal 8 percent and 9 percent, which match post-WWII U.S. averages. We assume there are three debt target states ($m = 3$): low, mid, and high. They are set to equal 60, 75, and 90 percent of output, which correspond to the House Republican's 2012 deficit reduction proposal (and the CBO's Baseline projection), the President's 2012 budget, and the CBO's alternative fiscal scenario.[6] The strength of the fiscal response to changes in the debt-to-output ratio, γ, is set to equal 0.3 so that a switch in the true debt target state takes the economy approximately 10 years to adjust to its new equilibrium in the full information case, which is consistent with Congress's planning horizon. It also guarantees a sufficient response by the fiscal authority to ensure stable long-run debt dynamics.

Figure 2a shows the speed of convergence of the debt-to-output percentage after a switch in the debt target state from mid (75) to low (60). The 2012 House proposal is provided for comparison. When $\gamma = 0.3$ the speed of convergence matches closely with the House's 10-year proposal. The implied 10-year target also matches closely. In the model, the long-run debt target is approached asymptotically, so there is a trade-off between matching the 10-year target and the speed of convergence. Thus, we further justify our choice of γ by at looking the short-run adjustment in the tax rate implied by the endogenous tax rule. A 3.75 percent short-run adjustment (figure 2b) seems consistent with current deficit reduction proposals, considering the 2012 American Taxpayer Relief Act changed marginal rates by similar amounts (the top marginal income tax rate rose by 4.6 and the payroll tax rose by 2 percentage points), but not enough to achieve the low debt target.

To obtain a historical sense of the degree of fiscal uncertainty, figure 3a plots each of the CBO's baseline projections of the debt-to-GDP percentage (dashed lines) with the actual percentage superimposed (solid line). It is not surprising that over long horizons these projections are uninformative. By law, the CBO's projections are based on current law, which is a poor predictor of long-run policy. For example, the projections in the early-mid 1990s did not initially account for the Balanced Budget Act of 1997 and in the early 2000s the projections did not initially account for the U.S. involvement in Iraq and Afghanistan or the Bush tax cuts in 2001. These policy changes suggest that the debt target switches over time and that households must form a probability distribution over these outcomes to set their expectations of future tax policies.

Figure 3b plots the distribution of the difference between the CBO's projection in any given

[6] Since 2007, the CBO has published its debt-to-GDP projections, given the assumptions in their baseline and alternative fiscal financing scenarios. Comparing these projections with the actual debt-to-GDP path from 2007-2012 reveals significant error. Thus, it is possible that households are conditioning on an even wider set of debt targets.

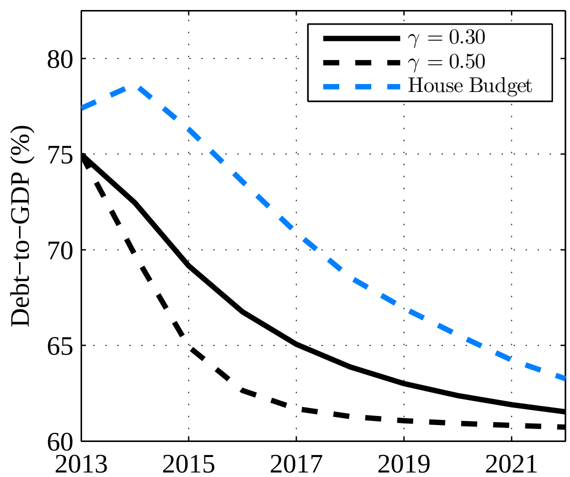

(a) Path to the long-run debt target

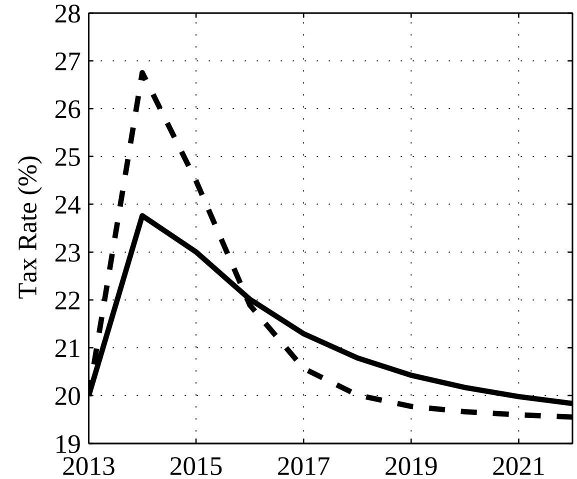

(b) Short-run income tax adjustment

Figure 2: The path of debt-to-GDP after a switch from the mid (75%) to low (60%) long-run debt target and the corresponding tax rates. Discretionary policy shocks are set to zero.

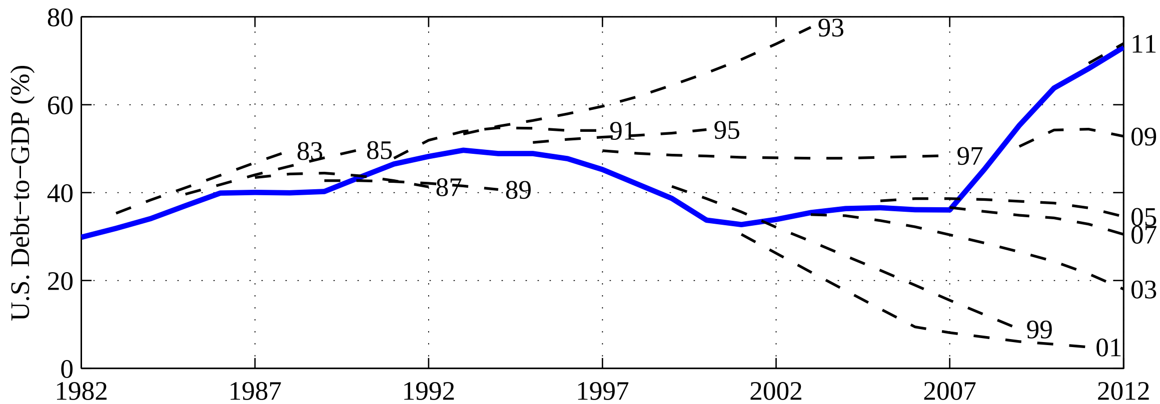

(a) CBO baseline projections of debt-to-GDP percentages (dashed lines) from Economic and Budget Outlook reports (1986-2010) and Five Year Budget Projection reports (1982-1984) compared to actual debt-to-GDP percentages (solid line).

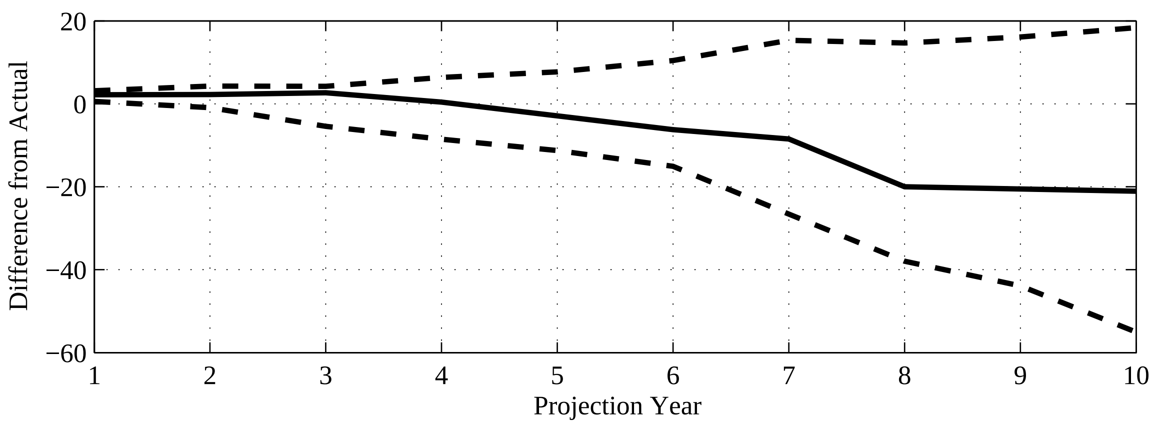

(b) Distribution of the difference between the CBO baseline projections of the debt-to-GDP percentage and the actual debt-to-GDP percentage, across projection years. The solid line is the median and the dashed lines are the 25/75 quantiles.

Figure 3: Comparison of CBO baseline projections of the debt-to-GDP percentages with the actual values

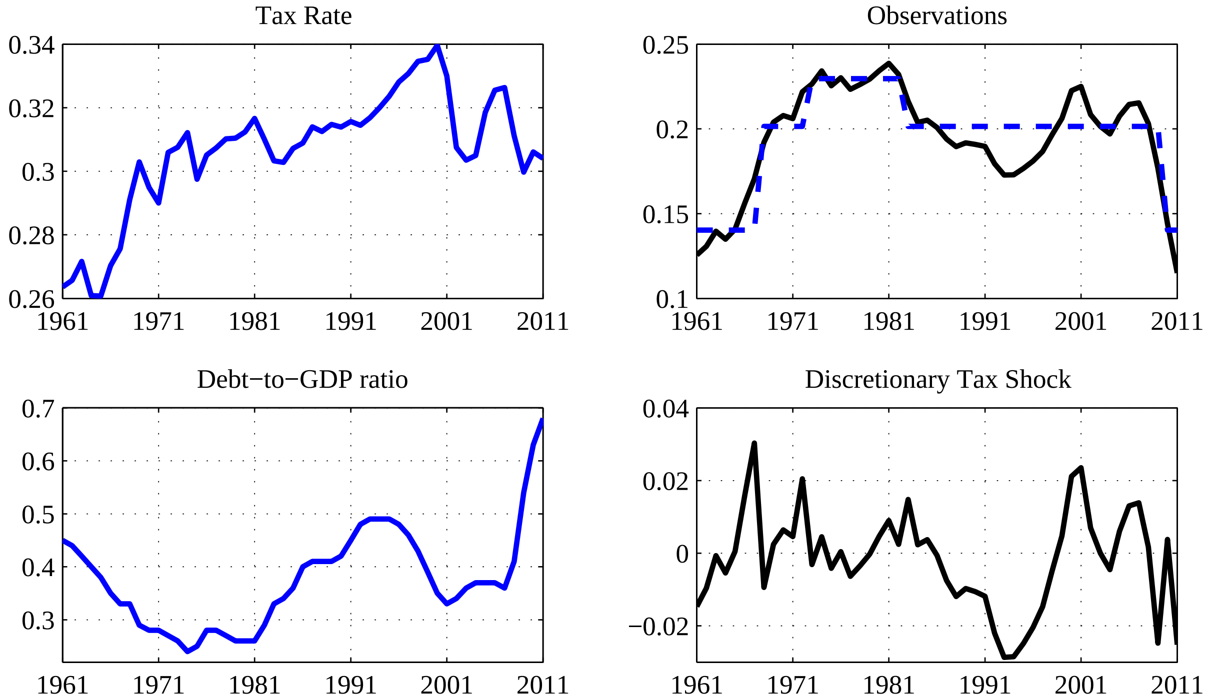

Figure 4: Tax revenue is Federal, State, and Local receipts (NIPA Table 3.2 Line 1 + Table 3.3 Line 1). GDP is nominal (Table 1.1.5 Line 1) since tax receipts are nominal. Debt-to-GDP is published in the Economic Report of the President. The debt-to-GDP ratio is from the Economic Report of the President. The solid line in the upper right panel are the household's observations calculated with the data and the endogenous tax rule where $\gamma = 0.30$. The implied state-dependent intercepts (dashed line) in the tax process are the 10/50/90 quantiles of the observations.

year and the actual debt-to-GDP percentage. Over shorter horizons, significant legislative changes are less likely to occur and, on average, the CBO's projections are more accurate. The accuracy of the CBO's projections over a 5-year time span suggest that changes in the debt target are infrequent.

Figure 4 gives a sense of how well the endogenous tax rule matches the transitions in the data. The left panel shows U.S. average tax rate and debt-to-GDP data from 1961 to 2011. The tax rate is calculated as the share of tax revenue to GDP from the NIPA tables. The upper right panel shows the household's observations, $x_t = \tau_t - \gamma b_{t-1}/y_{t-1}$, implied by the data over the sample. Applying the Gibbs sampler to this data under the assumptions that (10) is the data generating process, $\gamma = 0.30$, and there are three debt target states, yields a sampled sequence of states (with corresponding intercepts given by the 10/50/90 quantiles of the observations) shown with a dashed line. The implied discretionary tax shocks are shown in the bottom right panel and have a standard deviation of 0.013. The sampled average transition matrix and 68 percent credible interval are

$$P_{16} = \begin{bmatrix} 0.78 & 0.11 & 0.05 \\ 0.07 & 0.81 & 0.05 \\ 0.07 & 0.12 & 0.67 \end{bmatrix}, \quad \bar{P} = \begin{bmatrix} 0.81 & 0.13 & 0.07 \\ 0.08 & 0.84 & 0.08 \\ 0.10 & 0.19 & 0.71 \end{bmatrix}, \quad P_{84} = \begin{bmatrix} 0.83 & 0.15 & 0.08 \\ 0.10 & 0.87 & 0.10 \\ 0.12 & 0.22 & 0.79 \end{bmatrix}.$$

The point of this estimation procedure is to apply the same tool (i.e., the Gibbs sampler on the sequence of states and the transition matrix) to historical data that the household uses to form expectations about the future debt target and how it evolves. If we were to assume that $\bar{P}$ represents the true historical transition matrix, then it is reasonable to ask whether and how quickly the household is able to learn about this transition matrix over time. However, we choose a symmetric true transition matrix near the 68 percent credible interval of the sampled transition matrix to avoid a situation where, on average, low debt targets are more likely than high debt targets or vise versa in the baseline setup. Moreover, current policy does not indicate that one regime is more likely than another. Thus, we set the true transition matrix so that each debt target state has an expected duration of 5 years ($p_{ii} = 0.8$) and there are equal probabilities of switching to the other states ($p_{ij} = 0.1$ for $i \neq j$). Perhaps coincidentally, five years coincides with the persistence of current law (figure 3b) and the average number of years served by a president.

The standard deviation of the discretionary tax shock, σ_ε, is set equal to 0.02, which is roughly half the distance between the intercepts of the tax rule, (10). Under this calibration the household cannot easily distinguish between discretionary tax shocks and changes to the debt target state. For example, conditional on the mid debt target, there is at least a 35 percent chance that a discretionary tax shock yields a signal that is more than half way to the neighboring long-run targets.

5 RESULTS

We present three sets of results. The first set is based on a unique sequence of debt target states and discretionary tax shocks. It shows how uncertainty across the three information sets impacts the paths of key aggregate variables relative to the full information case and describes the sources of the differences. The second and third sets of results are based on $5{,}000$ Monte Carlo simulations of the model. These results show two measures of learning and the welfare consequences of the uncertainty across the information sets. We conclude by showing that the uncertainty created by the recent fiscal policy debate has slowed the recovery and led to welfare losses.

5.1 UNCERTAINTY CHANNELS

Figure 5 shows the differences between the paths of output, consumption, capital, and labor hours in the three limited information cases in percent deviations from the full information case (case 0). The purpose of this exercise is to illustrate how limited information changes the household's decisions relative to when they have full information. The most limited information case influences the business cycle by regularly moving output by 0.1 percent and capital by 0.3 percent (and sometimes more) relative to case 0 over the 100 year simulation. Thus, uncertainty about the long-run debt target has a similar effect on the business cycle as Fernandez-Villaverde et al. (2011) attribute to stochastic volatility shocks to fiscal policy.

The paths for case 1 (solid line) shown in figure 5 contain sequences of spikes that correspond to significant differences of the household's belief from the true long-run debt target, which are shaded over the path of output in figure 6. The top panel compares the household's belief, calculated as the average of s_t weighted by its likelihood, against the true state. In the middle panel, a dark-shaded/blue (light-shaded/red) region means the household's belief is higher (lower) than the true value by more than one standard deviation of the differences over the simulation. The direction is important because it has opposite implications for the household's expectation of the future tax rate. If they believe the debt target is higher (lower) than the truth, then they expect the tax rate will decrease (increase) to allow debt to rise (fall) toward its target. This is why case 1 output is generally higher (lower) than case 0 output when the belief is above (below) the truth. For example, the bottom panel shows a sequence of consecutive positive and relatively large tax shocks between years 40 and 46. These shocks cause the household to believe the debt target is low when in fact the middle target is the truth. The path of output in case 1 falls below case 0 due

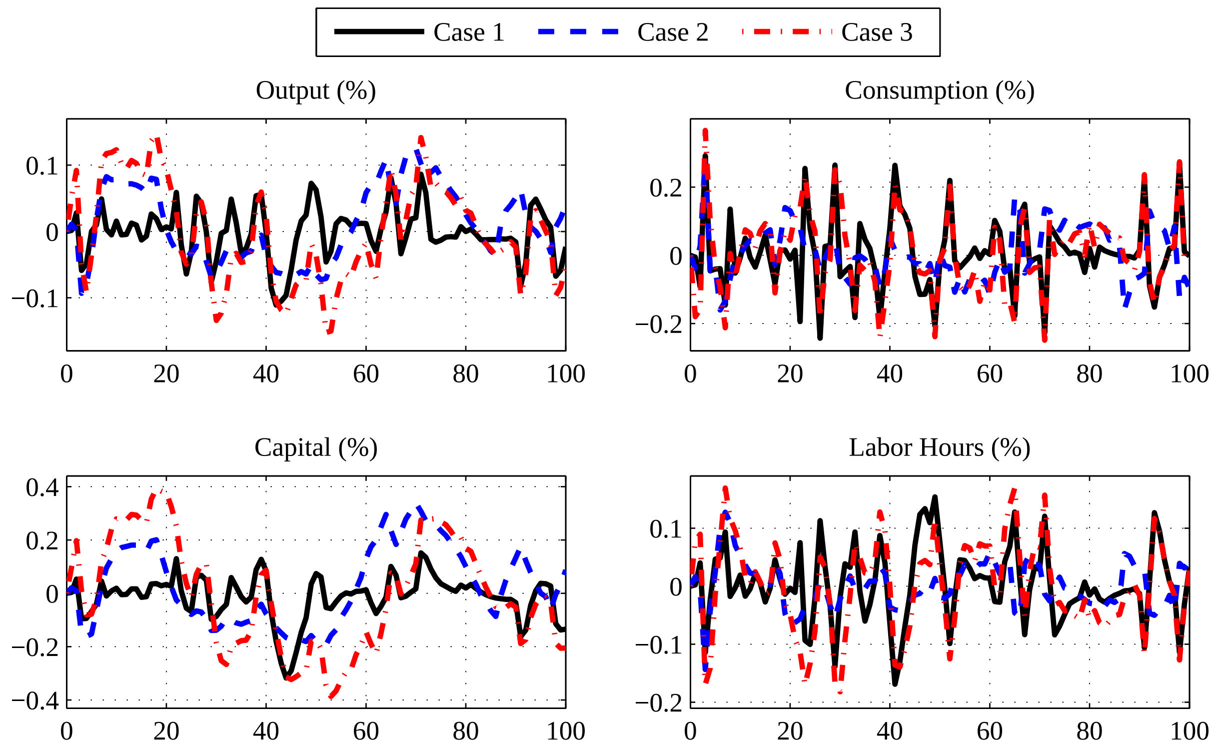

Figure 5: Simulated paths in percent deviations from case 0. The paths for case 1 (solid line), case 2 (dashed line), and case 3 (dash-dotted line) are based on identical sequences of debt target states and discretionary tax shocks.

to the belief that short-run tax rates will continue to be high to retire debt to a lower long-run level.

In case 1, the household rationally learns about the debt target state. They are aware that their current belief about the state may be wrong and account for the uncertainty by weighting each realization of the state with its likelihood. The deviations of case 1 from case 0 are a function of how far their belief is from the truth. Moreover, they are less responsive to discretionary tax shocks than in case 0. For example, it is unlikely that the household is sure that the debt target switched to a high (low) target. This means they put positive weight in expectation on the choices they would make conditional on the lower (higher) debt targets. Therefore, their choices are tempered by the uncertainty and rarely match the optimal consumption, labor, and investment decisions in case 0.

In case 2, the deviations from case 0 have different characteristics than case 1. Since the household knows the true transition matrix is time invariant, they use their current estimate of the transition matrix when forming expectations (i.e., they assume that it does not change). Initially, observations of different transitional events are rare, which creates bias in the estimates toward the few that have been observed. The influence of the prior helps to alleviate some of this bias, but at the same time smooths their initial beliefs. These facts cause more persistent deviations of the paths from case 0 than in case 1. However, as they gather more observations of the state, their estimate of the transition matrix converges to the truth. Since the transition matrix belief is the only difference between case 2 and case 0, the differences in the simulations disappear as time evolves.

In case 3, the unknown transition matrix creates smoother and more persistent deviations from case 0, which are interrupted by sequences of spikes when the belief of the state differs from the truth. There is also interaction between the unknown state and transition matrix. In general, estimates of the transition matrix in case 3 are farther away from the truth than in case 1. This

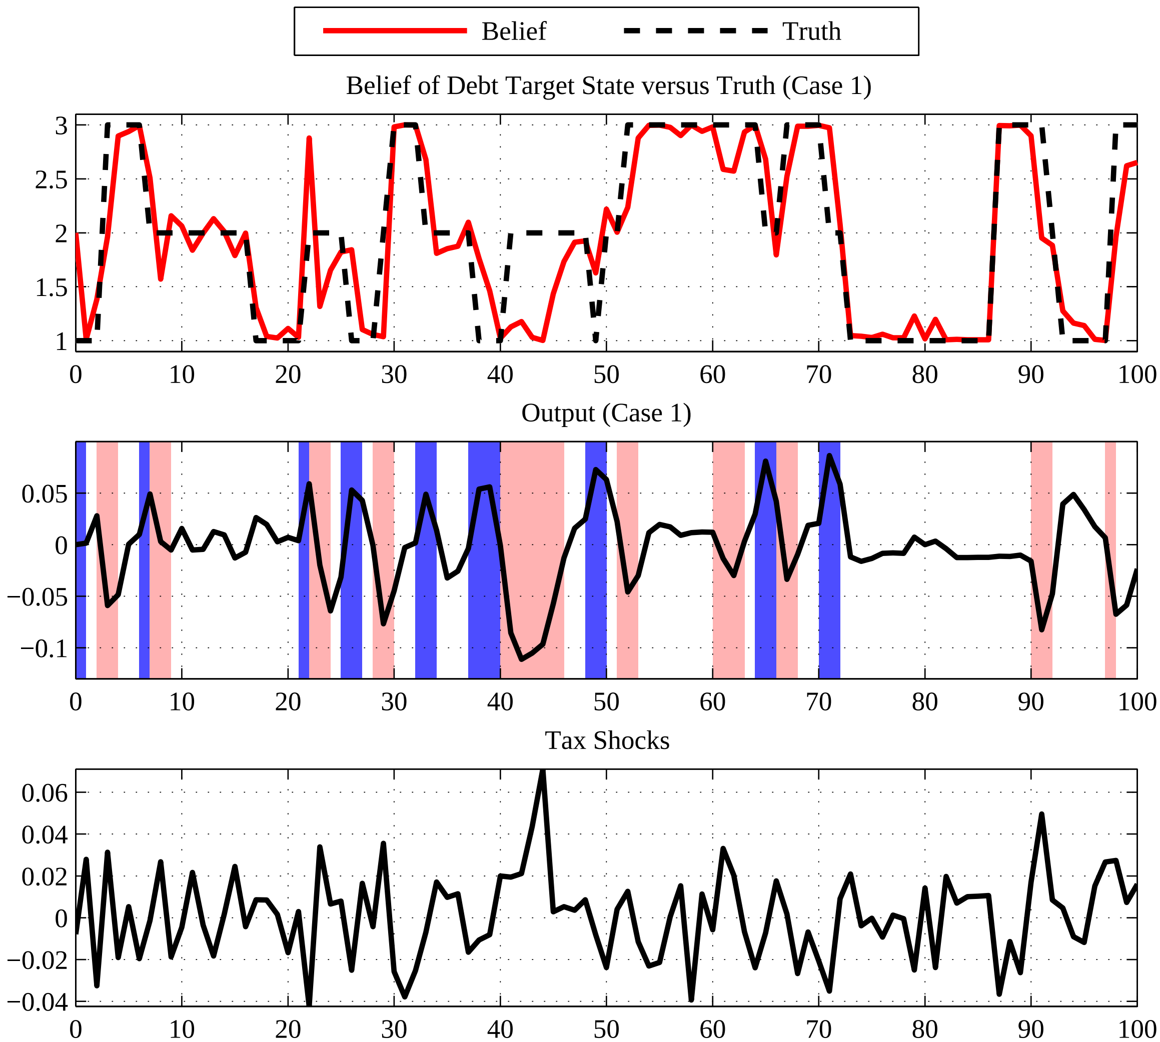

Figure 6: Sources of the deviations of Case 1 from Case 0. Top: Household's belief, calculated as the average of s_t weighted by its likelihood. Middle: Path of Case 1 output in percent deviations from Case 0 output. A dark-shaded/blue (light-shaded/red) region means the belief is higher (lower) than the true value by one standard deviation. Bottom: Tax shocks, which obscure the state and alter the household's beliefs.

is because $\{s_i\}_{i=0}^T$ is sampled sequentially with P, given the household's observations. When the state is known, only the transition matrix is sampled and the estimate converges to the truth. However, when the state is unknown, the uncertainty about the state adversely affects their ability to estimate the transition matrix. Regardless, as time evolves their estimates improve and converge toward the truth. The interaction of the unknown state with the estimate of the transition matrix causes the paths in case 3 to deviate from case 0 more than the paths in cases 1 and 2.

In summary, the unknown state causes high frequency deviations from the full information case, while the unknown transition matrix causes low frequency movements.

5.2 IMPULSE RESPONSES TO A TAX SHOCK

To get a better understanding for the dynamics in figures 5 and 6, we first plot the responses to a two standard deviation negative discretionary tax shock, conditional on full information, in figure 7. The responses are shown in deviations from the stochastic steady state and the debt target state is held constant at its middle realization. The negative shock reduces the tax rate in period 1. A lower tax rate increases the after-tax return to labor, which increases hours worked. This raises output, consumption, and investment. The larger tax base is not sufficient to offset the decrease in the tax rate and debt rises relative to GDP. Higher outstanding debt increases the tax rate above its steady state in period 2. This reduces labor hours, investment, and output. As time evolves, all variables slowly decay toward steady state.

Figure 8 identifies the effect of an incorrect belief in case 1. The impulse responses are shown as percent deviations of case 1 from case 0 to isolate the effects of the change in beliefs. The tax rate initially corresponds to the middle debt target and the debt-to-output ratio equals its corresponding stochastic steady state. Consider a scenario where the household incorrectly believes the debt target state rose in period 1 (solid line) due to a cut in the tax rate. Since the tax rate is based on the true debt target state and the past debt-to-output ratio, it does not deviate from case 0 in period 1. However, they increase their labor supply and substitute away from consumption in favor of investment, since the household expects lower future taxes. Also, they expect to be wealthier in the future and for consumption growth to increase, so the real interest rate increases. With a broader tax base, debt falls. In period 2, the household realizes their belief of the debt target was incorrect, and decreases their labor supply and investment in anticipation of higher future taxes. The higher real interest rate in period 1 causes debt to rise in period 2. In period 3, the tax rate increases in response to the increase in debt, which further reduces labor and investment. Thus, limited information about the debt target state causes the endogenous variables to differ from the case 0 paths, even after their beliefs are perfectly aligned with the truth.

5.3 MEASURES OF LEARNING We use two measures, computed from $5{,}000$ Monte Carlo simulations of the model, to assess how quickly and how effectively the household learns about the debt target state and how it evolves. The first is the distance of the estimated state likelihood vector from the true vector in cases 1 and 3, where the household does not know the sequence of long-run debt targets. The second is the distance between the estimated debt target transition matrix and the true matrix in cases 2 and 3. In both measures, we take the average of the distances (measured by the 2-norm) in each period across all simulations as a percentage of the maximum distance.

Figure 9a plots the distance of the estimated state likelihood vector from the true vector across time. In cases 1 and 3, the household knows the debt target transition matrix, but starts with an uninformative prior for the state likelihood since they begin the simulation without an observation to guide them. The household's estimate of the likelihood vector approaches the truth faster in case 1 than in case 3 since they know the true transition matrix. As time evolves, the household collects observations of the tax rate and the lagged debt-to-output ratio, which improves their ability to learn the true likelihood vector. In case 1, the true transition matrix is fed into the Hamilton filter, whereas in case 3 the matrix evolves from the uninformative prior, which alters the likelihood of each state in case 3 relative to case 1. Therefore, there is a uniformly smaller average distance between the belief and the truth in case 1 relative to case 3. By period 40, the household achieves very small improvements in their estimates, and most of the learning occurs in the first 10 years.

The household's ability to correctly infer the debt target state depends on the distance between the debt targets and the variance of the discretionary tax shock. Figure 10a compares the distance of the state likelihood belief from the truth when the volatility of the shock, σ_ε, decreases from 0.02 to 0.01. A higher variance increases the difficulty of the signal extraction problem by changing the distribution of the observations. Figure 10b shows a near normal distribution when $\sigma_\varepsilon = 0.02$

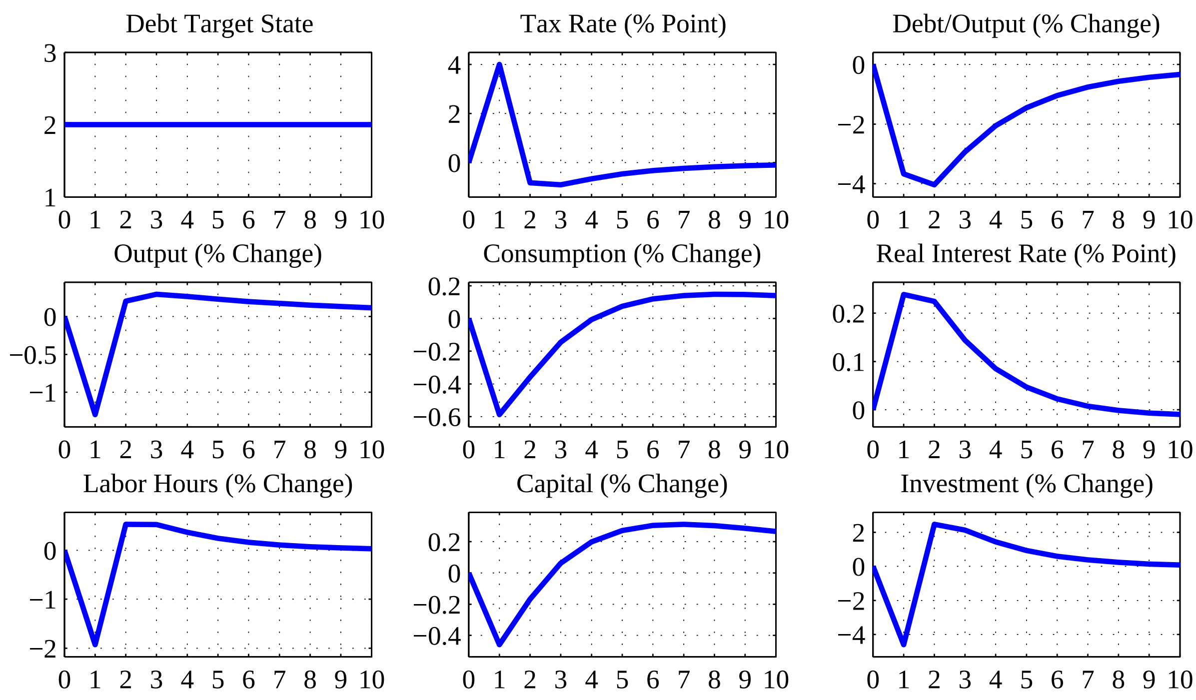

Figure 7: Case 0 impulse responses to a two standard deviation negative discretionary tax shock in percent deviations from the stochastic steady state.

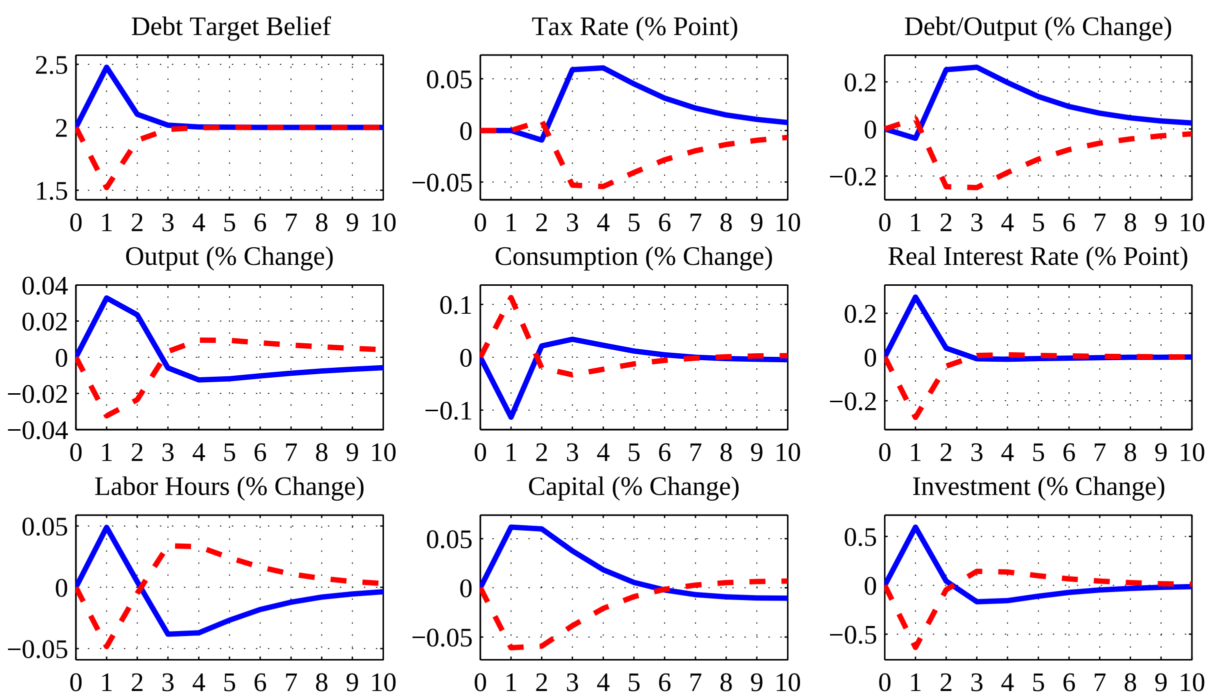

Figure 8: Case 1 impulse responses to a two standard deviation positive (dashed line) and negative (solid line) discretionary tax shock in percent deviations from full information (case 0). The true state is always the mid debt target.

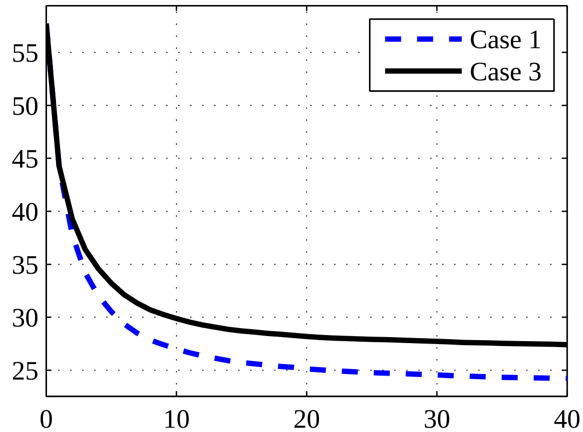

(a) Norm difference between true and estimated state likelihood

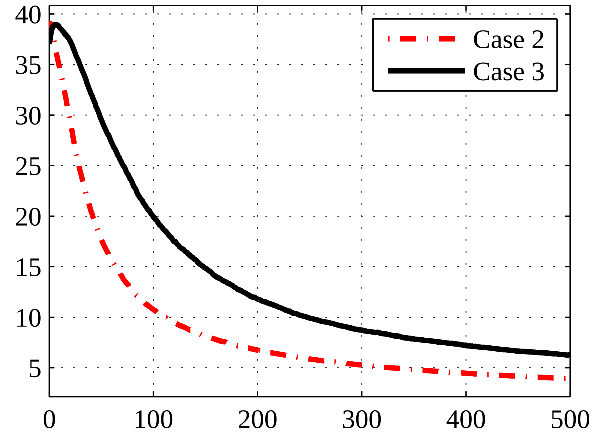

(b) Norm difference between true and estimated transition matrix

Figure 9: Alternative measures of learning under case 1 (dashed line), case 2 (dash-dotted line), and case 3 (solid line). The values are based on $5{,}000$ Monte Carlo simulations of the model.

and a clearly trimodal distribution when $\sigma_\varepsilon = 0.01$. The household's likelihood estimate nearly converges to the truth by period 40 in the low variance case, contrary to the high variance case.

Figure 9b plots the distance between the estimated and true transition matrices, which shows how long it takes for the household's beliefs to converge. The household learns the transition matrix much slower than they learn the unknown debt target state, because the household cannot take advantage of discrete shifts in the debt target, as they do when they learn the state, and they need a long history of observations to obtain accurate estimates. This stresses the importance of omitting the transition matrix from the household's information set. Case 2 converges more quickly than case 3 since the household knows the true debt target state. This is significant because it shows that the model converges to the rational expectations equilibrium.

The initial increase in case 3 (figure 9b, solid line) happens through period 5, the average duration of the current long-run debt target. Initial observations will reflect few, if any, switches in the state. This creates bias toward the corresponding diagonal probability in the transition matrix, which favors the states the household believes have occurred, while the other rows remain uninformative. Once the household infers a switch in the state, the estimate of the diagonal probability corresponding to staying in the previous state begins to correct itself. Thus, the estimates converge toward the truth as the household gathers more observations. However, case 3 converges slower than case 2 since the unknown state adversely affects the household's estimates.

5.4 WELFARE DISTRIBUTIONS This section quantifies the welfare consequences in each limited information case. We first describe how to calculate welfare and then show the distributions across time. The results are based on $5{,}000$ Monte Carlo simulations of the model.

Following the welfare cost measure of Schmitt-Grohe and Uribe (2007), we quantify the welfare costs of the uncertainty by thinking of the decisions made with limited information as alternative policies to the full information case. Thus, we solve for a λ^ℓ that satisfies

$$\mathbb{E}_t^\ell W(c_t(\mathbf{z}_{t-1}^\ell), n_t(\mathbf{z}_{t-1}^\ell)) = \mathbb{E}_t^0 W((1-\lambda^\ell)c_t(\mathbf{z}_{t-1}^\ell), n_t(\mathbf{z}_{t-1}^\ell)),$$

where $W(c_t(\mathbf{z}_{t-1}^\ell), n_t(\mathbf{z}_{t-1}^\ell)) \equiv \sum_{i=t}^{T-1} \beta^{i-t} u(c_i(\mathbf{z}_{i-1}^\ell), n_i(\mathbf{z}_{i-1}^\ell))$. T is the simulation length, so

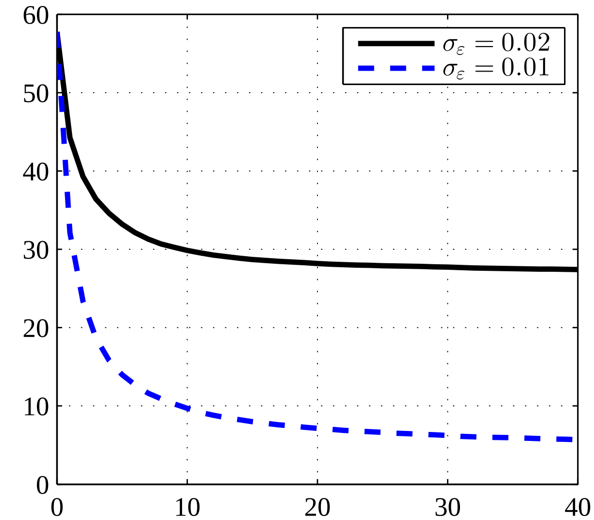

(a) Norm difference between true and estimated state likelihood

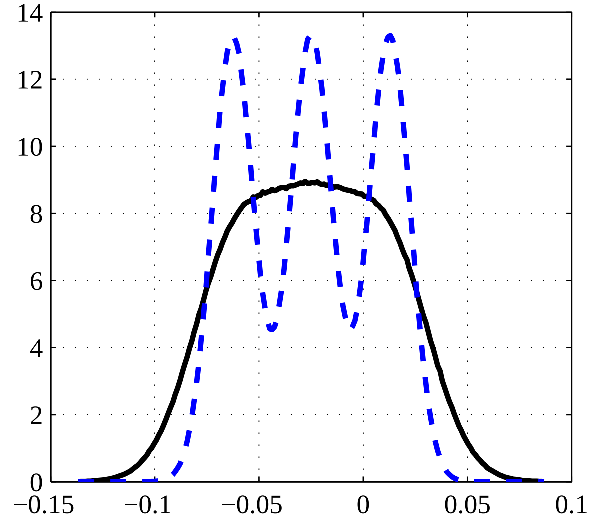

(b) Distribution of observations (probability density functions)

Figure 10: The effect of a change in the variance of discretionary tax shocks in case 3. Compares $\sigma_\varepsilon = 0.02$ (solid line) and $\sigma_\varepsilon = 0.01$ (dashed line). The values are based on $5{,}000$ Monte Carlo simulations of the model.

W is the time-t present value of remaining simulation utility. Since the calculation is based on a simulation, the length of the interval $[t, T]$ must ensure that the present value at t of utility near the end of the simulation is close to zero. We choose $T = t + 501$, noting that $\beta^{500} \approx 3 \times 10^{-9}$. With $\sigma = 1$, the additively separable specification of utility implies

$$\lambda^\ell = 1 - \exp\left\{\frac{1-\beta}{1-\beta^{T-t}}\left(\mathbb{E}_t^\ell W(c_t(\mathbf{z}_{t-1}^\ell), n_t(\mathbf{z}_{t-1}^\ell)) - \mathbb{E}_t^0 W(c_t(\mathbf{z}_{t-1}^\ell), n_t(\mathbf{z}_{t-1}^\ell))\right)\right\}.$$

A positive (negative) λ^ℓ represents a welfare cost (benefit) in limited information case ℓ relative to case 0. We initialize the welfare calculations in period 0 at the stochastic steady state and the calculations in period $t > 0$ at the limited information state vector from period $t - 1$.

Figure 11 compares the distribution of λ^ℓ across cases 1, 2, and 3 at the median (solid line) and one standard deviation bands (dashed lines). In case 1 (figure 11a), the household can be compensated for not knowing the current debt target state with ± 0.001 percent of full information consumption goods 68 percent of the time. In half of the simulations in a given period, uncertainty about the debt target state causes welfare benefits (a negative λ^1). As time evolves, the consumption goods required to compensate the household for the uncertainty does not diminish, despite the household's ability to learn about the state's likelihood vector very quickly.

How can an unknown debt target state lead to welfare gains? In case 0, the household knows the current state but not future states. This means the household must form expectations over the possible states to inform their optimal decisions. When the debt target state is unknown, the household's expectations are potentially misaligned with case 0, which may lead to welfare gains from upside risk. For example, if the household expects a lower debt target (higher taxes), then they invest less relative to case 0. If the debt target winds up higher than they expected, welfare is lower than case 0. With less investment, the household's after-tax returns are lower, which reduces future consumption, leisure, and welfare, despite initially higher consumption and leisure. However, if the debt target winds up lower than they expected, the household's lower level of

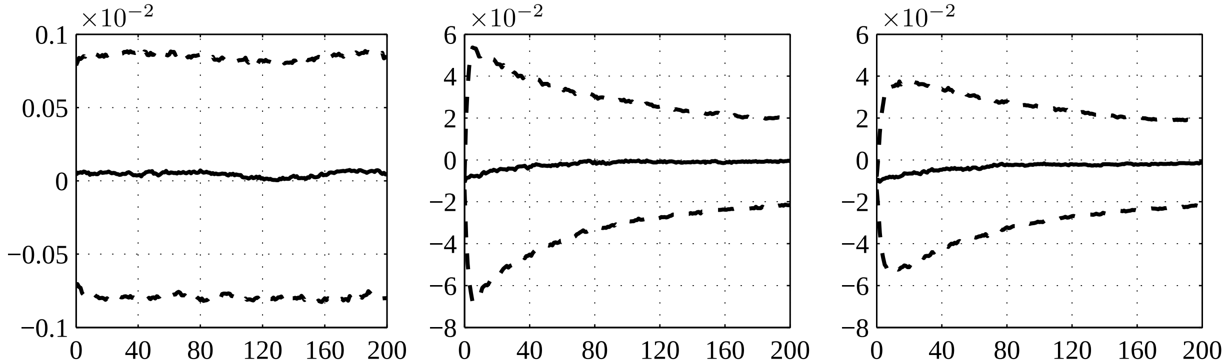

(a) Case 1 (b) Case 2 (c) Case 3

Figure 11: The distribution of welfare costs (λ^{ℓ}) as a percent of case 0 consumption goods forgone over the remaining simulation. The solid line is the median and the dashed lines are the $16/84$ percentile bands. A positive (negative) number represents a welfare cost (benefit). The values are based on $5{,}000$ Monte Carlo simulations with $\sigma_{\varepsilon} = 0.02$.

investment is better aligned with the realization of taxes, and welfare is higher than case 0. When the household's prior and true transition matrices are symmetric, these outcomes are equally likely.

Figure 11b plots welfare in case 2, which shows that the unknown transition probabilities have a very different effect than the unknown state. Three facts stand out relative to case 1: the welfare cost distribution remains fairly symmetric, the compensation is much larger in the tails, and the compensation diminishes as the household learns the transition matrix. The explanation for the symmetric distribution is similar to case 1, but the source of the welfare consequences is different. Given an uninformative prior on the transition matrix, the household initially expects all debt targets are visited with an equal probability. However, their initial observations can bias their estimates, which affects their decisions relative to case 0. More specifically, if the household observes a particular sequence of states that favors one target more than the others, then they will over-estimate the probability of switching to and remaining at that target, which misaligns their expectations with the truth. As an example, suppose a simulation remains in the high debt target state longer than the true expected duration (5 years). This causes the household to believe high debt targets (lower taxes) are more likely and to increase their investment relative to case 0. If the debt target remains high in the future, the household benefits since their after-tax returns will be higher than in case 0. However, if the debt target falls in the future, the household loses, since they face higher taxes. On average, half of the draws of the state will be shorter (longer) than the true expected duration, which leads to symmetric welfare benefits and costs.[7]

The household requires more consumption goods to compensate them for the unknown transition matrix than the unknown state for two reasons. First, the transition matrix is more complicated to learn. It is defined by six parameters and requires a large sample of observations to obtain an accurate estimate. Second, the household treats their estimate of the transition matrix as the truth in expectation. Their estimate changes across time, but in expectation they believe it is the truth and time invariant. Thus, their suboptimal decisions persist across time and the welfare consequences are more severe than in case 1. As the household gathers more observations, on average their estimates approach the truth, which reduces the variance of the welfare costs. Initially, the

[7]The full and limited information stochastic steady states differ, which causes slight bias in the welfare distribution. This is due to the anticipated utility approach to forming expectations.

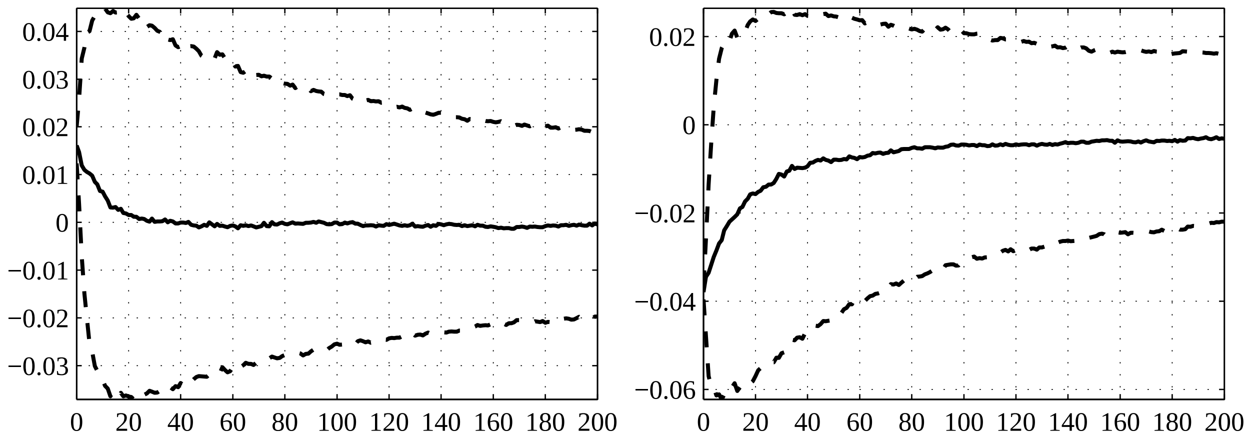

(a) High debt target visited more often (P_H)

(b) Low debt target visited more often (P_L)

Figure 12: A comparison of the distribution of welfare costs in case 3 (λ^3) in terms of the percent of case 0 consumption goods forgone over the remaining simulation when there is asymmetry in the true transition matrix. The solid line is the median and the dashed lines are the $16/84$ one standard deviation bands. A positive (negative) number represents a welfare cost (benefit). The values are based on $5{,}000$ Monte Carlo simulations.

household requires ± 0.06 percent of full information consumption goods to compensate them for an unknown transition matrix, but by period 200 the distribution tapers to ± 0.02 percent.

In case 3, there is initially less variance in the welfare distribution relative to case 2. When the household does not know the current debt target state or how it evolves, they evaluate expectations conditional on their estimated transition matrix and state likelihood vector. Since the household knows they do not know the current state, it tempers their choices relative to case 2 and shrinks the welfare consequences. Initially, the household requires ± 0.04 percent of full information consumption goods to compensate them for the uncertainty. This falls to ± 0.02 percent by period 200, which is similar to case 2 since the feedback effects between the unknown state and unknown transition probabilities decline as the household improves their estimates of the transition probabilities.

The fact that the welfare distributions are fairly symmetric is not surprising. The symmetry results from an equal probability of some simulations overrepresenting a particular state. This is because the true transition matrix and the household's prior have the same ergodic distribution (i.e., they expect to visit the three debt targets equally often). Figure 12 relaxes this assumption, while keeping the diagonal the same as the baseline calibration, with the alternatives

$$P_H = \begin{bmatrix} 0.80 & 0.10 & 0.10 \\ 0.05 & 0.80 & 0.15 \\ 0.05 & 0.15 & 0.80 \end{bmatrix} \quad \text{and} \quad P_L = \begin{bmatrix} 0.80 & 0.15 & 0.05 \\ 0.15 & 0.80 & 0.05 \\ 0.10 & 0.10 & 0.80 \end{bmatrix},$$

which have ergodic distributions $\Pi(P^L) = [0.4, 0.4, 0.2]$ and $\Pi(P^H) = [0.2, 0.4, 0.4]$. Note that with P_L (P_H) as the true transition matrix, the low (high) debt target is visited more often, but the average duration of all three states remains unchanged. When the true transition matrix has a different ergodic distribution than the household's prior, there is bias in the initial welfare distribution. Consider the case where the high debt target is visited more frequently (figure 12a). The household's prior belief is symmetric. Therefore, the household expects higher taxes and underinvests relative to the full information case. If the target winds up being even lower (even higher taxes), the household benefits, and if the target winds up being higher (lower taxes), the household loses.

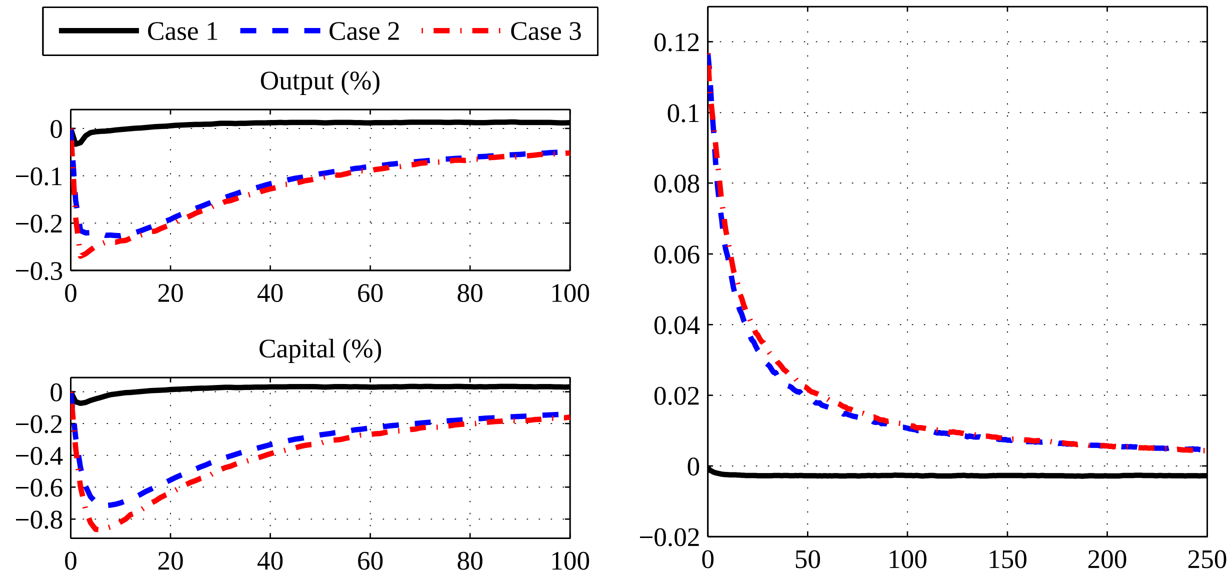

(a) Contractionary paths

(b) Welfare costs

Figure 13: Consequences of recent fiscal uncertainty. The values are given in percent deviations from case 0 and based on Monte Carlo simulations of the model where the true debt target is always high.

When the true transition matrix is symmetric, these forecast errors are equally likely. However, when the transition matrix is skewed so that high debt targets are more frequent, then it is less likely that the target winds up being lower, meaning it is less likely that the household benefits. Thus, on average, there are welfare losses at the beginning of the simulation. A similar argument explains why there are initially median welfare gains when the low debt target is visited more frequently (figure 12b). As the household learns the true transition matrix, the bias disappears.

5.5 UNCERTAINTY AND RECENT FISCAL POLICY

The previous section makes clear that limited information about the debt target creates two-sided risk. Depending on how the household's expectations are aligned with the realization of future states, limited information can result in gains or losses. These results may seem counterintuitive, given the popular belief that fiscal uncertainty leads to welfare losses. However, this section shows that the uncertainty created by the recent fiscal policy debate has depressed output and led to welfare losses. To understand why, consider the following scenario. The fiscal authority announces they will reduce deficits with an increase in taxes. The household believes the announcement is credible and places positive probability mass on higher short-run taxes. However, when it comes time to increase taxes, the fiscal authority reneges on its promise and keeps taxes low. The household is surprised to see taxes remain low and invested less than they would have if they knew taxes would stay low. This mistake leads to persistently lower output, which the economy recovers from as the household revises its expectations.

To see the quantitative effects of this policy, we consider an experiment where the true debt target state is always high (i.e., there is no uncertainty about future debt target states in case 0 since the high debt target state is perfectly absorbing), but the household faces limited information about the state and how it evolves. Figure 13a plots the paths of output and capital (in percent deviations from case 0), which are based on Monte Carlo simulations of this experiment. In case 1, learning about the high debt target state is not difficult, and the paths closely follow case 0. However,

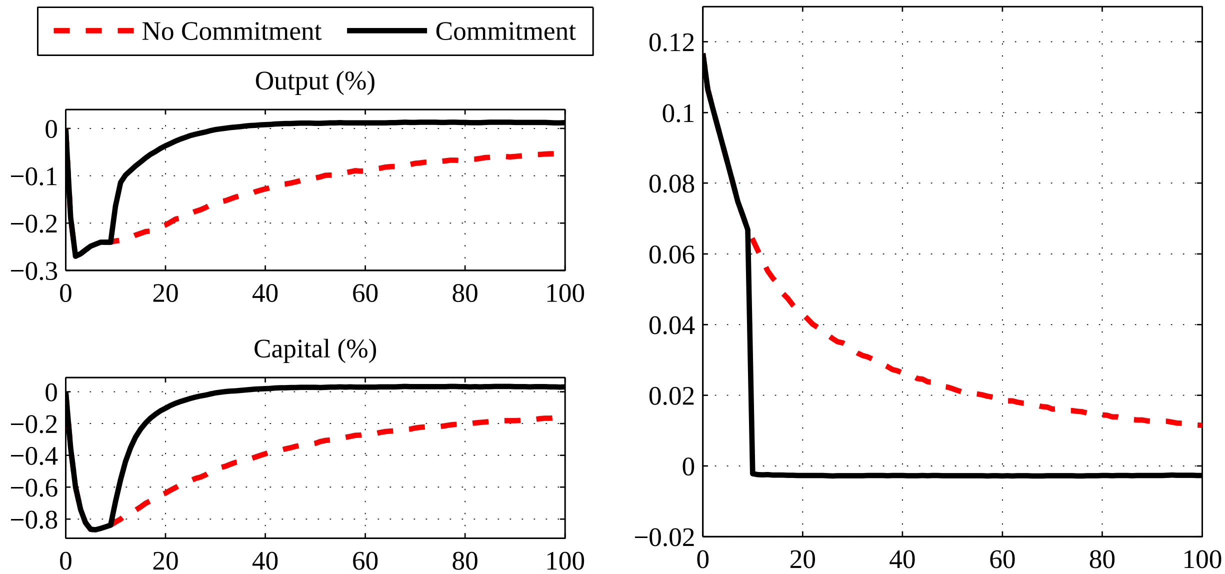

(a) Contractionary paths

(b) Welfare costs

Figure 14: Effect of a policy that credibly commits to the high debt target. The values are given in percent deviations from case 0 and based on Monte Carlo simulations of the model where the true debt target is always high.

in cases 2 and 3, learning about the transition matrix results in lower capital, consumption, and leisure. The household's initial belief of the transition matrix is symmetric so that realizations of the tax rate are strictly below expectations. Each period the household is surprised by low taxes. These surprises decline over time as the household's belief of the transition matrix approaches the truth. Initially, the surprises are large and the household invests much less than in case 0 due to higher expected taxes. Their belief of the transition matrix adjusts rapidly, which leads to higher investment and a build-up of the capital stock. Thus, capital accumulates to the level in case 0, at the expense of lower consumption, leisure, and welfare (figure 13b). This example indicates that the recent fiscal policy debate, which has led to heightened uncertainty, has slowed the recovery.

Figure 14 shows the effect of the fiscal authority deciding to credibly commit to the high debt target to guide the household's expectations after a prolonged period of uncertainty. The commitment in period 10 expedites the recovery from the uncertainty to full information output. The experiment is the same as figure 13, except beginning in period 10 the household's beliefs coincide with the true transition matrix (i.e., they know the high debt target is an absorbing state). Once the household's expectations match actual fiscal policy dynamics, output nearly makes a full recovery within 10 years. Therefore, the fiscal authority is able to effectively counteract the negative consequences of the household's limited information from the first 10 years of the simulation.

The speed of the recovery depends on how far the paths of output and capital initially deviate from full information, which is a function of the household's prior belief of the transition matrix. If the household is endowed with more information about the true transition matrix, then the initial impact of uncertainty is smaller and the recovery is quicker following the announcement of the fiscal authority's commitment. Regardless, uncertainty about the debt target always makes the household worse off if their expected tax rate under limited information is less consistent with the realization of taxes than their expected tax rate under full information. There are many scenarios when this is true. However, it is particularly acute if the household expects tax rates to strictly increase (decrease) when, in reality, they do not change. The scheduled sunset of the Bush tax cuts in 2010 is one recent example of such a policy since the tax cuts were extended for two years and then permanently extended for families (individuals) earning under $\$450{,}000$ ($\$400{,}000$) annually. People likely placed positive probability mass on the sunset provision (and little or no probability on further tax cuts) due to initiatives by the President and Congress to reduce deficits, but were surprised when the tax cuts were permanently extended. This scenario demonstrates that if households receive clearer signals about future policy, then they are less hesitant to invest and the economy quickly recovers, even if past signals were noisy.

6 CONCLUSION

The main contribution of this paper is to model the uncertainty surrounding the long-run debt target by assuming that households have limited information about both the current debt target state and the transition probabilities. A Bayesian household learns about the state and uses an anticipated utility approach to form expectations conditional on their belief about the transition matrix. We analyze the effects of the uncertainty on the household's decisions and quantify the welfare consequences across three limited information sets. Our key findings are as follows:

1. The household learns about the transition matrix much slower than they learn about the unknown debt target state. Moreover, incorrect inferences of the state adversely affect the household's estimates of the true transition matrix, creating feedback effects that alter the endogenous paths. Perhaps surprisingly, the variance of the welfare distribution is smaller when the debt target state is unknown because the household acts more cautiously.

2. Uncertainty creates persistent deviations from the full information paths. Discretionary tax shocks cause the household to incorrectly place probability mass on a change in the debt target state, which affects their consumption/saving and labor/leisure decisions. This changes the real interest rate and leads to persistent changes in the path of debt relative to output. Under the most limited information set (case 3), the paths of capital and output regularly deviate from the full information paths (case 0) by 0.3 and 0.1 percent.

3. Fiscal uncertainty imposes two-sided risk. When the household's prior and the true transition matrix are symmetric, there is an equal chance of welfare gains and losses with no welfare consequences in the median of the distribution. If the true transition matrix visits a high (low) debt target more often, then the household under-invests (over-invests) when tax rates are lower (higher) than expected, which shifts the welfare distribution toward losses (gains).

4. The uncertainty created by the recent policy debate has depressed output and led to welfare losses. Although a portion of the welfare distribution shows gains, these outcomes did not occur since the realizations of policy were inconsistent with people's expectations. When Congress provides greater clarity about the future direction of fiscal policy, even after a period of uncertainty, output and welfare increase and the economy quickly recovers.

We caution that our results represent a floor on the consequences of fiscal uncertainty, since uncertain debt targets is only one source. Uncertainty about government spending and entitlement programs represent additional layers of fiscal uncertainty that are important future research topics.

REFERENCES

ABBAS, S., N. BELHOCINE, A. ELGANAINY, , AND M. HORTON (2011): "Historical Patterns of Public Debt—Evidence From a New Database," *IMF Economic Review*, 59, 717–742.

AIZENMAN, J. AND N. P. MARION (1993): "Policy Uncertainty, Persistence and Growth," *Review of International Economics*, 1, 145–63.

ALBERT, J. H. AND S. CHIB (1993): "Bayes Inference via Gibbs Sampling of Autoregressive Time Series Subject to Markov Mean and Variance Shifts," *Journal of Business & Economic Statistics*, 11, pp. 1–15.

ALEXOPOULOS, M. AND J. COHEN (2009): "Uncertain Times, Uncertain Measures," University of Toronto Working Paper 352.

ANDOLFATTO, D. AND P. GOMME (2003): "Monetary Policy Regimes and Beliefs," *International Economic Review*, 44, 1–30.

BACHMANN, R. AND C. BAYER (2013): "'Wait-and-See' Business Cycles?" *Journal of Monetary Economics*, 60, 704–719.

BASU, S. AND B. BUNDICK (2012): "Uncertainty Shocks in a Model of Effective Demand," NBER Working Paper 18420.

BI, H., E. M. LEEPER, AND C. LEITH (2013): "Uncertain Fiscal Consolidations," *Economic Journal*, 123, F31–F63.

BIANCHI, F. AND L. MELOSI (2013): "Dormant Shocks and Fiscal Virtue," in *NBER Macroeconomics Annual 2013, Volume 28*, ed. by J. Parker and M. Woodford, University of Chicago Press.

BIZER, D. S. AND K. L. JUDD (1989): "Taxation and Uncertainty," *American Economic Review*, 79, 331–36.

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRYB (2012): "Really Uncertain Business Cycles," NBER Working Paper 18245.

BOHN, H. (1998): "The Behavior Of U.S. Public Debt And Deficits," *The Quarterly Journal of Economics*, 113, 949–963.

BORN, B. AND J. PFEIFER (2013): "Policy Risk and the Business Cycle," CESifo Working Paper 4336.

CHETTY, R., A. GUREN, D. S. MANOLI, AND A. WEBER (2011): "Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," NBER Working Paper 16729.

CHUGH, S. K. (2012): "Firm Risk and Leverage Based Business Cycles," Manuscript, Boston College.

COGLEY, T. AND T. J. SARGENT (2008): "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making," *International Economic Review*, 49, 185–221.

COMMITTEE FOR A RESPONSIBLE FEDERAL BUDGET (2011): "Analyzing the President's New Budget Framework," CRFB, Washington, D.C.

DAVIG, T. (2004): "Regime-Switching Debt and Taxation," *Journal of Monetary Economics*, 51, 837–859.

DAVIG, T. AND A. T. FOERSTER (2013): "Uncertainty and Fiscal Cliffs," Manuscript, Federal Reserve Bank of Kansas City.

DAVIG, T. AND E. M. LEEPER (2006): "Fluctuating Macro Policies and the Fiscal Theory," in *NBER Macroeconomics Annual 2006, Volume 21*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, MIT Press, Cambridge, 247–298.

——— (2011): "Temporarily Unstable Government Debt and Inflation," *IMF Economic Review*, 59, 233–270.

DAVIG, T., E. M. LEEPER, AND T. B. WALKER (2010): "'Unfunded Liabilities' and Uncertain Fiscal Financing," *Journal of Monetary Economics, Carnegie-Rochester Conference Series on Public Policy*, 57, 600–619.

——— (2011): "Inflation and the Fiscal Limit," *European Economic Review, Special Issue on Monetary and Fiscal Interactions in Times of Economic Stress*, 55, 31–47.

DOTSEY, M. (1990): "The Economic Effects of Production Taxes in a Stochastic Growth Model," *American Economic Review*, 80, 1168–82.

EVANS, G. W. AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*, Hoboken, NJ: Princeton University Press.

——— (2009): "Learning and Macroeconomics," *Annual Review of Economics*, 1, 421–451.

FERNANDEZ-VILLAVERDE, J., P. GUERRON-QUINTANA, K. KUESTER, AND J. RUBIO-RAMIREZ (2011): "Fiscal Volatility Shocks and Economic Activity," Federal Reserve Bank of Philadelphia Working Paper, 11-32.

GIANNITSAROU, C. (2006): "Supply-side Reforms and Learning Dynamics," *Journal of Monetary Economics*, 53, 291–309.

HAMILTON, J. D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–84.

JOHANNSEN, B. K. (2013): "When are the Effects of Fiscal Policy Uncertainty Large?" Manuscript, Northwestern University.

JUDD, K. L. (1998): *Numerical Methods in Economics*, Cambridge, Massachusetts: The MIT Press.

KREPS, D. (1998): "Anticipated Utility and Dynamic Choice," in *Frontiers of Research in Economic Theory*, ed. by D. Jacobs, E. Kalai, and M. Kamien, Cambridge: Cambridge University Press, 242–74.

LEEPER, E. M. AND T. ZHA (2003): "Modest Policy Interventions," *Journal of Monetary Economics*, 50, 1673–1700.

POOLE, K. T. AND H. ROSENTHAL (2012): "The Polarization of the Congressional Parties," URL: http://voteview.com/political_polarization.asp.

RICHTER, A. (2012): "The Fiscal Limit and Non-Ricardian Consumers," Manuscript, Auburn University.

RICHTER, A. W., N. A. THROCKMORTON, AND T. B. WALKER (2013): "Accuracy, Speed and Robustness of Policy Function Iteration," *Forthcoming, Computational Economics*.

SCHMITT-GROHE, S. AND M. URIBE (2007): "Optimal Simple and Implementable Monetary and Fiscal Rules," *Journal of Monetary Economics*, 54, 1702–1725.

SCHORFHEIDE, F. (2005): "Learning and Monetary Policy Shifts," *Review of Economic Dynamics*, 8, 392–419.

A TECHNICAL APPENDIX

A.1 DISCRETIZATION METHOD (3-STATE MARKOV CHAIN) Several Markov switching models contain only two states. If it is a hidden Markov model, then a single probability, call it q_t, describes the household's belief about the state in period t. The household believes with probability q_t that they are in state 1 and probability $1-q_t$ that they are in state 2. Thus, if the state is unknown at t and there are only two realizations, then discretizing q on the interval $[0,1]$ in the state-space is straightforward. Policy function iteration is used to obtain an approximate solution to the hidden Markov model by including q in the state-space with an appropriate rule to update q to form expectations. However, if there are m states in the Markov chain, then the restriction $\sum_i \mathbf{q}_t(i) = 1$, for $i \in \{1,\ldots,m\}$ where $\mathbf{q}_t(i) = \Pr(s_t = i|\mathbf{x}^t)$, imposes a constraint on the discretized state space. For example, if each $\mathbf{q}_t(i)$ is simply discretized on the interval $[0,1]$, then the solution procedure will attempt to find approximating policy functions on regions of the state-space that violate the restriction on the likelihood (e.g., $\mathbf{q}_t(1) = \mathbf{q}_t(2) = \mathbf{q}_t(3) = 1$ or $\sum_i \mathbf{q}_t(i) = 3$). Therefore, an alternative approach to discretizing the state space is necessary.

Suppose that $m = 3$, as in our model. The restriction on the likelihood in period t is $\mathbf{q}_t(1) + \mathbf{q}_t(2) + \mathbf{q}_t(3) = 1$. Note that given two of the three probabilities, the third is defined; geometrically, the restriction is a formula for a plane. The additional restrictions that $0 \leq \mathbf{q}_t(i) \leq 1$ make it an equilateral triangle with vertices at the unit vectors $[1,0,0]$, $[0,1,0]$, and $[0,0,1]$. Therefore, there exists a one-to-one mapping (bijection), $g : \mathbb{R}^3 \to \mathbb{R}^2$, that satisfies the above restrictions (i.e., one that defines the triangle in a 2-dimensional space). If $g(\mathbf{q}) = \boldsymbol{\theta}$, with $\mathbf{q}$ and $\boldsymbol{\theta}$ as row vectors, then the mapping g is a projection,

$$g(\mathbf{q}) = (\mathbf{q} - \mathbf{o})\mathbf{B} = \boldsymbol{\theta},$$

where $\mathbf{o}$ defines the origin of the projection in $\mathbb{R}^3$, $\mathbf{B}$ is a orthonormal basis of the original plane, and $\sum_i \mathbf{q}_t(i) = 1$. To derive $\mathbf{B}$, note that $\tilde{\mathbf{b}}_1 = [0,1,-1]$ and $\tilde{\mathbf{b}}_2 = [1,0,-1]$ form a basis for the original plane. We use the Gram-Schmidt process to obtain orthogonal vectors $\mathbf{b}_1$ and $\mathbf{b}_2$, where

$$\mathbf{b}_1 = \tilde{\mathbf{b}}_1 = [0,1,-1], \qquad \mathbf{b}_2 = \tilde{\mathbf{b}}_2 - \mathrm{proj}_{\mathbf{b}_1}\left(\tilde{\mathbf{b}}_2\right) = [1,-1/2,-1/2],$$

so that $\mathbf{B} \equiv [\mathbf{b}_1^T/||\mathbf{b}_1||, \mathbf{b}_2^T/||\mathbf{b}_2||]$ is an orthonormal basis. To see the mapping clearly, define $\mathbf{o} \equiv [1,0,0]$ and write the mapping as a system of two linear equations,

$$\theta_1 = (\mathbf{q}_t(1) - 1)b_{11} + \mathbf{q}_t(2)b_{21} + \mathbf{q}_t(3)b_{31}$$
$$\theta_2 = (\mathbf{q}_t(1) - 1)b_{12} + \mathbf{q}_t(2)b_{22} + \mathbf{q}_t(3)b_{32}.$$

Imposing $\sum_i \theta_i = 1$ and simplifying the system implies

$$\theta_1 = \mathbf{q}_t(2)(b_{21} - b_{11}) + \mathbf{q}_t(3)(b_{31} - b_{11})$$
$$\theta_2 = \mathbf{q}_t(2)(b_{22} - b_{12}) + \mathbf{q}_t(3)(b_{32} - b_{12}).$$

Thus, given an appropriate choice of $\mathbf{B}$, there is a one-to-one mapping between $\boldsymbol{\theta}$ and $\mathbf{q}$.

To implement the discretization, $\mathbf{q}_t(2)$ and $\mathbf{q}_t(3)$ are first discretized on the interval $[0,1]$, with $\mathbf{q}_t(1)$ implied by the likelihood restriction. Permutations that violate $\mathbf{q}_t(1) \geq 0$ are discarded. The remaining coordinates are mapped into $\boldsymbol{\theta} \in \mathbb{R}^2$, which is included in the state-space and has the practical benefit of reducing the state-space by one dimension. The policy functions are approximated with ordinary least squares on a complete monomial basis of order 2 to avoid complications with ordering the discretized coordinates in a meaningful way.

A.2 HAMILTON FILTER In cases 1 and 3, the household observes the tax rate and lagged debt-to-output ratio, but not the decomposition of the observation between the state-dependent parameters and the discretionary shock. This means the household does not know the current debt target state, s. With each additional observation, the household uses a Hamilton (1989) filter to update their belief about the probability distribution of s. This section outlines the filter, which is an application of Bayes theorem. The filter takes as input a vector of conditional probabilities, which are in the state-space of the household's optimization problem, given by,

$$\Pr[s_{t-1} = i | \mathbf{x}^{t-1}],$$

where $i \in \{1, 2, 3\}$ and $\mathbf{x}^{t-1} \equiv \{x_0, \ldots, x_{t-1}\}$, and updates them with the most recent observation,

$$x_t \equiv \tau_t - \gamma b_{t-1}/y_{t-1}.$$

The filter outputs the updated conditional probabilities, given by,

$$\Pr[s_t = j | \mathbf{x}^t],$$

where $j \in \{1, 2, 3\}$, and the conditional likelihood of x_t, $f(x_t|\mathbf{x}^{t-1})$. To apply the filter, note that the probability density of x_t is normally distributed so that

$$f(x_t | s_t = j, s_{t-1} = i, \mathbf{x}^{t-1}) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\varepsilon_t^2}{2\sigma^2} \right),$$

where $\varepsilon_t = x_t - (\bar{\tau}(s_t) - \gamma \overline{by}(s_t))$. The filter's output is obtained with the following sequence of calculations for all i, j:

1. Calculate the joint probability of $(s_t = j, s_{t-1} = i)$ conditional on $\mathbf{x}^{t-1}$,
$$\Pr[s_t = j, s_{t-1} = i | \mathbf{x}^{t-1}] = \Pr[s_t = j | s_{t-1} = i] \Pr[s_{t-1} = i | \mathbf{x}^{t-1}].$$
2. Calculate the joint conditional density-distribution of x_t and $(s_t = j, s_{t-1} = i)$,
$$f(x_t, s_t = j, s_{t-1} = i | \mathbf{x}^{t-1}) = f(x_t | s_t = j, s_{t-1} = i, \mathbf{x}^{t-1}) \Pr[s_t = j, s_{t-1} = i | \mathbf{x}^{t-1}].$$
3. Calculate the likelihood of x_t conditional on $\mathbf{x}^{t-1}$,
$$f(x_t | \mathbf{x}^{t-1}) = \sum_{j=1}^{m} \sum_{i=1}^{m} f(x_t, s_t = j, s_{t-1} = i | \mathbf{x}^{t-1}).$$
4. Calculate the joint probabilities of $(s_t = j, s_{t-1} = i)$ conditional on $\mathbf{x}^t$,
$$\Pr[s_t = j, s_{t-1} = i | \mathbf{x}^t] = \frac{f(x_t, s_t = j, s_{t-1} = i | \mathbf{x}^{t-1})}{f(x_t | \mathbf{x}^{t-1})}.$$
5. Calculate the output by summing the joint probabilities over the realizations s_{t-1},
$$\Pr[s_t = j | \mathbf{x}^t] = \sum_{i=1}^{m} \Pr[s_t = j, s_{t-1} = i | \mathbf{x}^t].$$

The household's expectations contain future policy variables, which are interpolated at the updated probability distribution for s_t.

A.3 IMPORTANCE SAMPLER In case 2, the household knows the debt target state, but does not know the probabilities of the transition matrix. However, they know there is a time invariant transition matrix, P, with stationary distribution, $\Pi(P)$. This section outlines the importance sampler, which the household uses to estimate the transition probabilities. The likelihood of observing $\mathbf{s}^T \equiv \{s_t\}_{t=0}^T$ is

$$f(\mathbf{s}^T|\Pi, P) = \left(\prod_{j=1}^{3} \Pi_j(P)^{\mathbf{1}_j}\right)\left(\prod_{i=1}^{3}\prod_{j=1}^{3} p_{ij}^{m_{ij}^o}\right) = \prod_{j=1}^{3} \Pi_j(P)^{\mathbf{1}_j} \prod_{j=1}^{3} p_{1j}^{m_{1j}^o} \prod_{j=1}^{3} p_{2j}^{m_{2j}^o} \prod_{j=1}^{3} p_{3j}^{m_{3j}^o}, \quad (14)$$

where Π_j is the jth element of the stationary distribution, $\mathbf{1}_j \in \{0, 1\}$ is an indicator for whether state j is occupied by the household at time 0, p_{ij} is the probability corresponding to row i and column j of the transition matrix, and m_{ij}^o is the number of observed transitions from states i to j. Following Geweke (2005), the conjugate prior follows a Dirichlet distribution given by

$$f(P) = \left[\prod_{i=1}^{3} \Gamma\left(\sum_{j=1}^{3} a_{ij}\right) \Big/ \prod_{i=1}^{3}\prod_{j=1}^{3} \Gamma(a_{ij})\right]\left(\prod_{i=1}^{3}\prod_{j=1}^{3} p_{ij}^{a_{ij}-1}\right), \quad (15)$$

where $a_{ij} > 0$ are the shaping parameters of the prior distribution, and Γ is the gamma function. The posterior density is given by the product of the likelihood function, (14), and the prior density, (15). Thus, dropping the constants of proportionality, the posterior density is

$$f(P|\mathbf{s}^T) \propto \left(\prod_{j=1}^{3} \Pi_j(P)^{\mathbf{1}_j}\right)\left(\prod_{i=1}^{3}\prod_{j=1}^{3} p_{ij}^{a_{ij}+m_{ij}^o-1}\right). \quad (16)$$

The number of observed transitions, which are calculated after inferring the current state, s_t, determine the shaping parameters of the posterior distribution. Thus, as time unfolds the households's estimates of the probabilities converge toward the true probabilities.

The posterior distribution, (16), does not correspond to any standard density function that we can sample from directly. However, the second component is a product of two independent Dirichlet probability density functions. Consequently, we utilize the following importance sampling algorithm to sample from this distribution. First, sample L draws from a Dirichlet distribution with parameters $\mathbf{a}_i + \mathbf{m}_i$ where $\mathbf{a}_i = [a_{i1} \quad a_{i2} \quad a_{i3}]$ and $\mathbf{m}_i = [m_{i1} \quad m_{i2} \quad m_{i3}]$ for $i = \{1, 2, 3\}$. To accomplish this, we follow the procedure outlined in Gelman et al. (2004, Appendix A). Draws are made from a Gamma distribution with parameters $((\mathbf{a}_i + \mathbf{m}_i)/2, 2)$, and then weighted by the sum of the draws corresponding to each state, i. That is, using the draws, x_{ij}^ℓ, from a Gamma distribution, a draw from a Dirichlet distribution is $\theta_{ij}^\ell = x_{ij}^\ell / \sum_{i=1}^3 x_{ij}^\ell$. Second, the draws from the Dirichlet distribution are weighted by the coefficient of the posterior distribution, $w_\ell \equiv \prod_{j=1}^3 \Pi_j(P_t^\ell)^{\mathbf{1}_j}$ for all $\ell \in \{1, \ldots, L\}$, and then divided by the sum of the weights. Formally, the weighting procedure produces an estimate of the transition probabilities (draw from the posterior distribution),

$$\hat{p}_{ij} = \frac{\sum_{\ell=1}^{L} w_\ell \theta_{ij}^\ell}{\sum_{\ell=1}^{L} w_\ell}.$$

Alternatively, we could have drawn from the posterior distribution by applying the independence Metropolis-Hastings algorithm or acceptance sampling. Since we use a representative agent model, the efficiency gains from these alternative sampling algorithms is negligible.

A.4 GIBBS SAMPLER In case 3, when both the current debt target state and transition probabilities are hidden, the household employs a Gibbs sampler to estimate the transition probabilities (i.e., they construct a Markov chain such that the limiting distribution of the chain is the joint distribution of interest). The household still uses the Hamilton filter to infer the current, end-of-sample, state. This is because the Gibbs sampler efficiently relies on the lagged and future states to draw a sequence of states each iteration, and so it performs relatively poorly at the end-of-sample where the future is unobserved. Furthermore, it cannot replace the Hamilton filter in the optimization problem since it requires the entire sample to estimate the probability distribution of s, which is computationally prohibitive.

Albert and Chib (1993) outline a Gibbs sampler to estimate an entire autoregressive 2-state hidden Markov model with r lags, but since we assume the household understands certain aspects of fiscal policy—the speed of expansion/consolidation, the long-run debt targets and supporting tax rates, and the volatility of *i.i.d.* discretionary tax shocks—the following Gibbs sampler only needs to iterate on the sequence of states and the transition probabilities. We chose a Gibbs sampler rather than maximum likelihood estimation of the Hamilton filter since imposing a prior distribution on the transition matrix yields reasonable estimates with short samples (i.e., it keeps the transition probabilities bounded away from from 0 and 1) and it does not require a nonlinear solver.

In period T, the household samples a sequence of states, $\mathbf{s}^T$, conditional on the 3-state transition matrix, P, and the observations, $\mathbf{x}^T$, sequentially. The following steps outline the sampler:

1. Initialize the Gibbs chain by sampling a sequence of states, $\mathbf{s}^T = \{s_1, \ldots, s_T\}$, from the prior transition matrix, P, using a uniform random number generator.

2. For $t \in \{1, \ldots, T\}$ and $j \in \{1, 2, 3\}$, sample s_t conditional on $\mathbf{x}^T$ and the neighboring states from the previous draw using the following rules:

   - If $t = 1$, then $f(s_1|\mathbf{x}^T, \mathbf{s}_{-1}) \propto \Pi_j(P) p_{jk} f(x_1|s_1)$, where $s_2 = k$.
   - If $1 < t < T$, then $f(s_t|\mathbf{x}^T, \mathbf{s}_{-t}) \propto p_{ij} p_{jk} f(x_t|s_t)$, where $s_{t-1} = i$ and $s_{t+1} = k$.
   - If $t = T$, then $f(s_T|\mathbf{x}^T, \mathbf{s}_{-T}) \propto \Pi_j(P) p_{ij} f(x_T|s_T)$, where $s_{T-1} = i$.

   s_t and x_t are particular realizations at time t, $\Pi_j(P)$ is the jth element of the stationary distribution of P, p_{ij} is the probability corresponding to row i and column j of P, and $\mathbf{s}_{-t} \equiv \mathbf{s}^T \setminus \{s_t\}$. The conditional probability density function, $f(x_t|s_t) = \exp\{-\varepsilon_t^2/(2\sigma^2)\}/\sqrt{2\pi\sigma^2}$, where $\varepsilon_t = x_t - (\bar{\tau}(s_t) - \gamma\overline{by}(s_t)))$ is the discretionary *i.i.d.* tax shock. s_t is drawn from each of these 3-element vectors of conditional probabilities with a uniform random number generator after normalizing by the sum of the elements.

3. Use the importance sampler in appendix A.3 to draw a new transition matrix, P.

4. Repeat steps 2 and 3 N times.

The household uses the average P over the last half of the Gibbs chain as their estimate of the transition matrix, $\hat{P}$, in period T.

B ADDITIONAL RESULTS (NOT FOR INTENDED PUBLICATION)

B.1 MEASURES OF UNCERTAINTY The results in this section are based on the same sequence of debt target states and discretionary tax shocks that generate the results in section 5.1.

Figure 15 shows the likelihood value associated with each debt target state in case 3. The square, x, and diamond markers correspond to the likelihood of being in the state with a high, middle, or low debt target. The solid line corresponds to the highest likelihood value across all states in a given period. This is a first pass at measuring the uncertainty about the state. The lower its value, the more uncertain the household is that the current state is the one given by the highest likelihood value. The drawback with this measure of uncertainty is that it ignores the possibility of the other two states having non-zero likelihood values.

Figure 16 presents an alternative measure of the uncertainty. Each period the Hamilton filter outputs the probability of each state conditional on the household's observations, $\Pr[s_t = i|\mathbf{x}^t]$. Absolute uncertainty occurs when the household assigns each state equal probability, $\Pr[s_t = i|\mathbf{x}^t] = 1/3$ for $i \in \{1,2,3\}$. In this case, we want the uncertainty index to equal 1. Absolute certainty occurs when the household assigns $\Pr[s_t = i|\mathbf{x}^t] = 1$ for some i. In this case, we want the uncertainty index to equal 0. The norm difference between the likelihood vector with absolute certainty and absolute uncertainty is $\text{norm}_{\max} = (\sum_{i=1}^3 (\mathbf{1}_i - 1/3)^2)^{1/2}$, where $\mathbf{1}_i$ is an indicator that equals 1 if the household is certain that state i is the truth and 0 otherwise. Thus, we adopt the following uncertainty index,

$$\frac{\text{norm}_{\max} - \sqrt{\sum_{i=1}^3 (\Pr[s_t = i|\mathbf{x}^t] - 1/3)^2}}{\text{norm}_{\max}},$$

which ranges from 0 (absolute certainty) to 1 (absolute uncertainty) as desired. The intervals where the difference in the Household's belief from the truth exceeds one standard deviation are superimposed on this uncertainty index (shaded regions). Notice the average level of uncertainty is greater in case 3 than in case 1. Periods of high uncertainty generally correspond to intervals where beliefs are farther from the truth, which alters the predicted business cycles relative to the full information case. However, it is possible the household could be very confident in their inference, when in fact they are incorrect, if there is a sequence of unusually large discretionary tax shocks (e.g. periods 40-46). In this case, they are simply mistaking a discretionary tax shock for a change in the long-run debt target, which causes them to be overconfident in their inference of the state.

B.2 FORECAST ERRORS Figure 17 shows the average variance of the consumption forecast errors as a percentage of steady state consumption across $5{,}000$ Monte Carlo simulations of the model, which provides an additional measure of how limited information affects the household. Let $\eta_t^c \equiv c_t - \mathbb{E}_{t-1}(c_t)$. Then the variance of the forecast errors through period t in a particular simulation is

$$\text{var}(\{\eta_i^c\}_{i=2}^t) = \frac{1}{t-2}\sum_{i=2}^t \left(\eta_i^c - \sum_{j=2}^t \eta_j^c/(t-1)\right)^2,$$

which is averaged across all simulations for each period t. The variance of the forecast errors is constant for cases 0 and 1, but the variance is smaller for case 1 even though the household has less information. This is because the household behaves more cautiously, which results in lower volatility of the endogenous variables and more predictable consumption. In cases 2 and 3, the

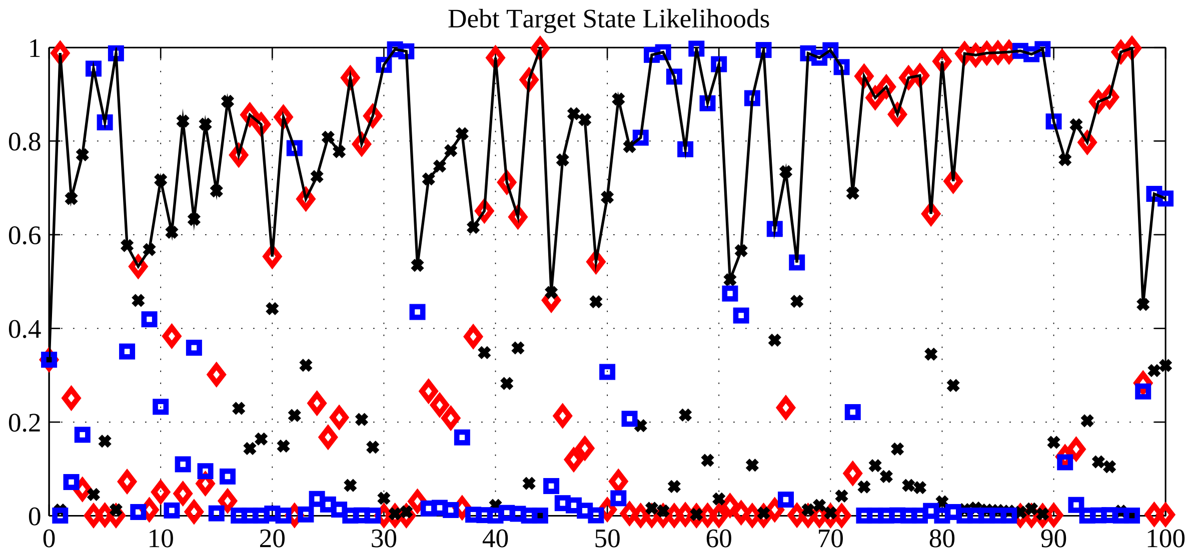

Figure 15: The likelihood of each state in case 3. The solid line is the highest likelihood across all 3 states. The square, x, and diamond markers correspond to the states with a high, middle, or low debt target intercept.

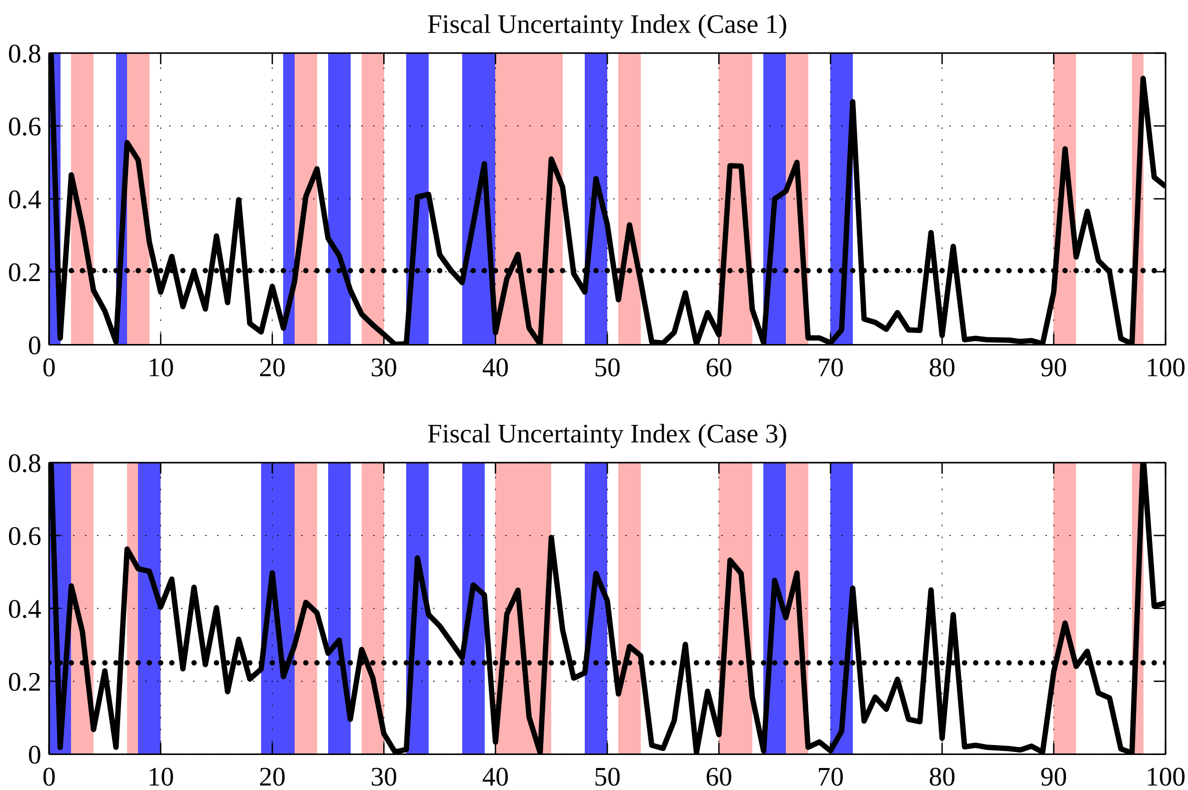

Figure 16: The uncertainty index is the norm difference between the vector of state likelihoods and absolute uncertainty (i.e. $(\sum_{i=1}^{3}(P[s_t = i|\mathbf{x}^t]-1/3)^2)^{1/2}$), as a percentage of the maximum possible norm difference. Thus the uncertainty index ranges from 0 (absolute certainty) to 1 (absolute uncertainty). The dashed line is the average value of the index. The top and bottom panels are based on identical sequences of debt target states and discretionary tax shocks.

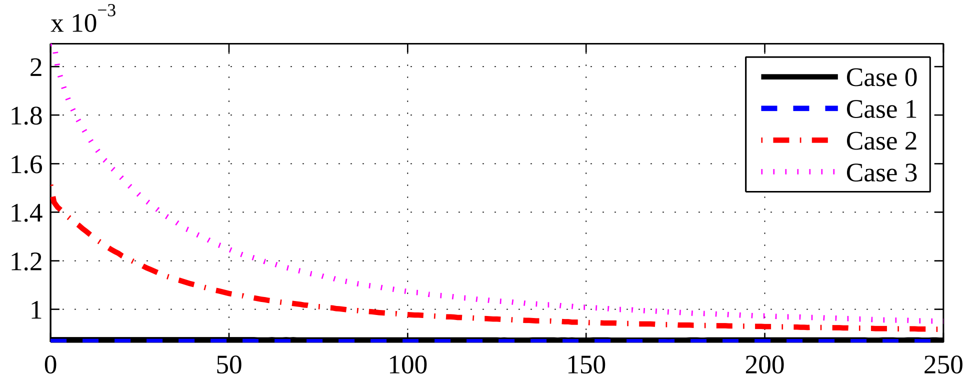

Figure 17: The average variance of the consumption forecast errors, as a percentage of deterministic steady state consumption. The values are based on $5{,}000$ Monte Carlo simulations of the model.

unknown transition matrix significantly increases the size of the forecast errors. As the household learns about the true transition matrix, the variance of the forecast errors decline. In case 3, more volatile forecast errors results from the unknown state reducing the household's ability to learn the true transition matrix. After 250 periods, the estimated transition matrix is close to the truth in both cases 2 and 3 and so the variance of the forecast errors approaches the variance in cases 0 and 1.

B.3 VARIANCE COMPARISON Figure 18 shows the median and one standard deviation bands of the distribution of welfare costs in case 3 (λ^3) in terms of the percent of case 0 consumption goods forgone over the remaining simulation and compares these distributions when the variance of the tax shock is high and low. Under either variance, welfare volatility increases at the beginning of the simulation, since the amount of uncertainty is high. The bands (dashed lines) taper toward zero as the influence of the uninformative prior in the Gibbs sampler diminishes and the household's estimate of the transition matrix becomes more accurate. Notice that the tails of the distribution are wider at the beginning of the simulation under low variance. Once again, this provides evidence that the household acts cautiously since the unknown state smooths their expectations.

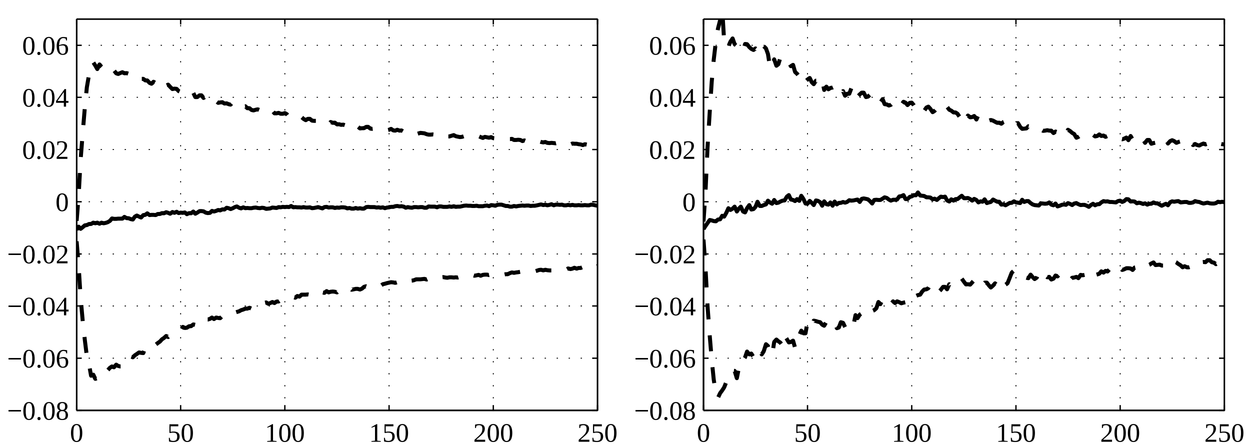

(a) High variance ($\sigma_\varepsilon = 0.02$)

(b) Low variance ($\sigma_\varepsilon = 0.01$)

Figure 18: The distribution of welfare costs in case 3 in terms of the percent of case 0 consumption goods forgone over the remaining simulation. The solid line is the median and the dashed lines are the $16/84$ percentile bands. A positive (negative) number represents a welfare cost (benefit). The values are based on $5{,}000$ Monte Carlo simulations.