Fiscal Policy with Heterogeneous Agents and Incomplete Markets

Jonathan Heathcote*

Duke University

July 28, 2001

Abstract

I undertake a quantitative investigation into the short run effects of changes in the timing of taxes for model economies in which heterogeneous households face a borrowing constraint. A combination of the distortionary effects of non-lump-sum taxation and the liquidity effects arising from the asset market structure are found to imply large real effects from tax changes. For example, a temporary proportional income tax increase in the benchmark model economy reduces aggregate consumption by around 29 cents for every additional dollar of tax revenue raised. The consumption of low wealth households who are close to the borrowing constraint is most sensitive to the current tax rate. While there are many such households, richer households account for a disproportionately large fraction of aggregate income and consumption. Thus the distortionary effects of proportional taxation are quantitatively more important at the aggregate level than the effects associated with incompleteness of asset markets.

Keywords: Ricardian equivalence; Heterogeneous agents; Borrowing constraints; Fiscal policy

JEL classification: E62; H24; H31; H63

---

*Correspondence to: Department of Economics, Duke University, Durham, NC 27708. E-mail: heathcote@econ.duke.edu. I am greatly indebted to my advisor Andrew Atkeson for all his advice and suggestions. I also thank Lee Ohanian, Jose-Victor Rios-Rull, Richard Rogerson, Randall Wright, Nicholas Souleles, Stephen Zeldes, Morten Ravn, Narayana Kocherlakota, Anton Braun, Paul Klein, and Kjetil Storesletten.

1. Introduction

The Ricardian insight, revisited by Barro (1974), is that with dynastic households, lump-sum taxes, and perfect capital markets, changes in the timing of taxes should not affect households' optimal consumption decisions. Thus the Ricardian theory predicts an equivalence in terms of prices and allocations between any time paths for taxes that imply the same total present value for tax revenue. In contrast to this theoretical result, a large amount of empirical work suggests that the timing of taxes *does* matter. For example, Bernheim (1987) argues that "virtually all [aggregate consumption function] studies indicate that every dollar of deficits stimulates between \$0.20 and \$0.50 of current consumer spending". In the hope of reconciling the apparent gap between the Ricardian view and the empirical evidence, various authors have explored quantitative theoretical models in which one or more of the conditions for Ricardian equivalence are not satisfied.

Ricardian equivalence may fail if a tax cut reduces the tax burden on the current generation at the expense of future generations. However, even the extreme assumption of zero inter-generational altruism does not appear capable of rationalizing the magnitude of observed deviations from Ricardian equivalence. For example, assuming taxes rise to stabilize debt after a one year tax cut, Poterba and Summers (1987) find a marginal propensity to consume (MPC) of about 0.06. Hubbard and Judd (1986) examine five-year deficits repaid over periods of either 10 or 20 years and find MPCs of between 0.03 and 0.05. The explanation for these small numbers is straightforward. Households *do* treat the fraction of a tax cut that will be paid for by the next generation as an addition to net wealth. However, households want to smooth any increase in consumption over the remainder of their lifetimes, and average life expectancy is long relative to the duration of the tax cuts considered.

This paper focuses on the effects of tax changes when at least one of the remaining two assumptions underpinning the Ricardian result is not satisfied. First, when taxes are distortionary, changes in the timing of distorting taxes affect the optimal inter-temporal allocation of labor effort, consumption and investment (see, for example, Auerbach and Kotlikoff 1987, Trostel 1993, Braun 1994, and McGrattan 1994). Second, if asset market imperfections are such that some households in the economy would like to borrow but cannot find credit, then these households will adjust consumption in response to temporary tax changes (see Hubbard and Judd 1986, Altig and Davis 1989, Daniel 1993, and Feldstein 1988).

I describe a dynamic general equilibrium model in which infinitely-lived heterogeneous households receive idiosyncratic shocks to labor efficiency which cannot be insured. Households can reduce the sensitivity of consumption to income changes by accumulating precautionary holdings of a single asset. However, if household asset holdings ever reach zero then further dis-saving is prohibited; households face a "no-borrowing" constraint. Since households differ in their productivity histories, the model generates an endogenous cross-sectional distribution of asset holdings. The government in the model economy finances constant government consumption by issuing debt and levying proportional income taxes. The tax rate is stochastic, so households face risk at the aggregate as well as at the idiosyncratic level.

This modelling framework can address the two channels identified above as having the potential to generate quantitatively large deviations from Ricardian equivalence. Changes in the income tax rate temporarily alter the returns to saving and to working, encouraging inter-temporal substitution in consumption and labor supply. The intuition for why the no-borrowing constraint generates real effects from tax changes is straightforward. Households that are unfortunate enough to have both very low asset holdings and low current income would like to borrow against future income to increase consumption. They are unable to do so because of the no-borrowing constraint. If the government cuts taxes, such households can now increase consumption by the extent to which the tax cut raises disposable income.

I assume that households have rational expectations, and at each point in time assign the correct probability to any future sequence for tax rates; in this sense there are no surprise shocks to the tax rate. Because real government spending is assumed constant, the model appropriately isolates the effect of changes in tax rates that are not accompanied by simultaneous adjustments to government consumption or lump-sum transfers. The analysis is conducted in a general equilibrium framework to incorporate the effects of aggregate tax shocks on the real wage and the real interest rate.

Every household in the economy chooses optimally how much to adjust consumption in response to a tax change. However, tax changes will likely have larger effects on aggregate variables the greater the fraction of households that are wealth-poor and thus potentially borrowed-constrained. I therefore specify the process for labor productivity so that the model endogenously generates a distribution for asset holdings resembling that in the United States. At the same time, the productivity process is restricted to be consistent with empirical estimates of earnings risk from the PSID. The process for taxes in the model is such that the share of aggregate output paid in taxes has the same persistence and variance as in the post-war United States, while the ratio of debt to GDP remains bounded.

In addition to the benchmark model which features both distortionary taxation and incomplete markets, I also consider various alternative model economies. First I study an economy with lump-sum taxes, since this is a natural framework for assessing the importance of the asset market structure alone as a propagation mechanism. To isolate the distortionary effects of stochastic taxes I then investigate a complete markets environment. Finally, I describe an economy designed to capture some features of the life-cycle, as a way to gauge potential interaction between life-cycle savings dynamics and borrowing constraints.

The main finding of the paper is that a combination of distortionary taxation and capital market imperfections can give rise to quantitatively important departures from Ricardian equivalence. For example, when the asset holding distribution resembles that in the United States, an income tax cut from a rate of 34.3 percent to a rate of 31.8 percent is associated with an immediate increase in aggregate consumption of 28.7 cents for each dollar of tax revenue lost.[1] Consideration of the variations on the benchmark model described above suggest that most of this effect is attributable to the distortionary nature of the tax system rather than the presence of the no-borrowing constraint. For example, a similar tax cut in a model with complete markets generates a 23.0 cent increase in consumption per dollar of revenue lost.

The rest of the paper is organized as follows. In the next section I review the empirical evidence on the response of aggregate consumption to tax changes, and the evidence on the importance of liquidity constraints at the household level. Section 3 contains a description of the model economies, along with a discussion of the choices for parameter values and the numerical solution methods. Section 4 discusses simulation results and assesses the relative importance of the distortions associated with proportional taxes versus the liquidity effects associated with a borrowing limit. Section 5 concludes.

[1] The long run implications of debt accumulation in my economy are the same as those in Aiyagari and McGrattan 1998, who find that increasing the steady state level of debt crowds out aggregate capital, reducing per capita consumption. The welfare cost of this reduction in the average *level* of household consumption is offset by a reduction in the average *volaility* of household consumption, since a higher real interest rate makes assets less costly to hold and therefore more effective in smoothing individual consumption. Woodford (1990) examines similar questions in a more stylized model.

2. Review of the Literature

2.1. Evidence on the response of consumption to tax changes

There is a large and rather inconclusive literature that tests for Ricardian equivalence (RE) by estimating consumption functions or Euler equations on aggregate time series (see the surveys in Bernheim 1987 and Seater 1993). One explanation for the lack of consensus is the problem of endogeneity. Cardia (1997) illustrates how the coefficient on the current budget deficit in an estimated consumption function (in which both output and the budget deficit are treated as independent variables) may be uninformative regarding the validity of RE if output responds immediately to tax changes. A second potential problem is that if current tax changes imply expected future government expenditure changes, then consumption might respond even if RE is true. As a third example, even if RE is false, consumption might only respond to *unanticipated* tax changes; this is a central implication of the permanent income / life cycle hypothesis (PILCH) model.

Given these difficulties, several authors have looked at various interesting natural experiments in which households saw large and reasonably well-understood changes in their disposable income. Various studies of the 1968 surtax and the 1975 rebate find quite large changes in aggregate consumption from these explicitly temporary tax changes. Modigliani and Steindal (1977) use large scale econometric models and estimate a marginal propensity to consume (MPC) over two quarters out of the 1975 rebate of between 0.3 and 0.58. Blinder (1981) examines both tax changes using a model based on the permanent income hypothesis and estimates a MPC of 0.16 over a quarter. Poterba (1988), using an Euler equation-based estimation, reports a MPC of between 0.13 and 0.27 within a month.[2] Wilcox (1989) finds large effects on consumption from the sequence of increases in social security benefits since 1965, even though these increases were always announced at least six weeks in advance.

Studies based on micro data have typically found even larger consumption responses to policy-induced income changes. Looking at the pre-announced Reagan tax cuts and using data from the Consumer Expenditure Survey (CEX), Souleles (2001) estimates a very large MPC for non-durables of between 0.6 and 0.9. Parker (1999), also using the CEX, estimates a MPC for nondurable goods of 0.20 for income changes associated with predictable changes in social security tax

[2] Poterba also finds that consumption did not appear to respond significantly to the passage of five large tax bills (including the 1968 and 1975 changes), even though it did respond when these tax changes were eventually implemented. The finding that aggregate consumption responds to predictable tax changes is in principle consistent with optimal forward-looking behavior if some.households are borrowing constrained.

with-holding. Souleles (1999) finds the MPC out of predictable income tax refunds to be between 0.35 and 0.6 within a quarter. Finally, Shapiro and Slemrod (1995) report that 43 percent of survey respondents planned to save most of the extra disposable income they would get from the 1992 reduction in the standard rate of with-holding for income taxes.

This apparent sensitivity of U.S. consumption to predictable changes in taxes or transfers is often attributed to the presence of liquidity constraints. It is therefore important to ask what other evidence (in addition to the response of consumption to tax changes) supports the view that borrowing constraints affect a large fraction of the population.

2.2. Evidence of the importance of borrowing constraints

Work on panel data indicates that some households in the U.S. *do* face liquidity constraints. Moreover, there appears to be a high correlation between the households that are liquidity constrained and those that have very little wealth. Zeldes (1989) works with the PSID and identifies the wealth-poorest and richest households in the sample. He rejects a permanent income hypothesis-based Euler equation for the poor, estimates a positive missing multiplier (suggesting they face a binding borrowing constraint), and finds that they exhibit excess consumption growth. Further cross-sectional evidence consistent with the presence of borrowing constraints is that households with low asset holdings appear to consume too little and have too little debt (see Hayashi 1985, and Cox and Jappelli 1993).[3]

Using data from the 1983 Survey of Consumer Finance, Jappelli (1990) finds that 12.5 percent of households report having requests for credit rejected, while a further 6.5 percent do not apply because they expected credit to refused. Thus, according to this measure, 19 percent of the U.S. population was liquidity constrained on at least one date in the year or two prior to the survey. Jappelli also finds that 74.1 percent of those households whose net worth is less than 15 percent of their disposable income are liquidity constrained (compared to 8.3 percent of those with greater net worth), suggesting that wealth-poor households are much more susceptible to finding themselves in the position of wishing to borrow but being unable to find credit.

The borrowing limit in the model described below is set equal to zero. This may be thought of either as an *ad hoc* borrowing limit or as the appropriate

[3] Souleles 1999 finds that on receipt of tax refunds, the nondurable consumption of those with low asset holdings rises much more than that of the rich. However, neither Souleles 2001 nor Parker 1999 find much evidence of a link between low asset holdings and excess sensitivity of consumption to predictable changes in income.

endogenous constraint for an economy in which there is no punishment for default. Because both the model and the empirical evidence imply a close connection between the characteristics of having low wealth and being unable to borrow, it is important to know how many wealth-poor households there are in the United States. Diaz-Gimenez, Quadrini and Rios-Rull (1997) report that in 1992 the poorest 40 percent of households held only 1.35 percent of total wealth, that approximately 3.4 percent of households had zero wealth, and that another 3.5 percent had negative wealth (suggesting that these households were able to take out imperfectly collateralized loans). Overall, these numbers suggest that a large fraction of the population may be at or near to their borrowing limit.[4]

3. The Models

The benchmark model is based upon the economies described in Aiyagari (1994) and Aiyagari and McGrattan (1998). A large (measure 1) number of households are *ex ante* identical and infinitely lived. They maximize expected discounted utility from consumption and from leisure. In aggregate, household savings decisions determine the evolution of the aggregate capital stock, which in turn determines aggregate output and the return to saving.

Because one goal of the paper is to assess the potential importance of liquidity constraints, I assume that households face idiosyncratic labor productivity shocks, and that markets which in principle could allow complete insurance against this risk do not exist. Instead there is a single risk-free savings instrument which enables households to partially self-insure by accumulating precautionary asset holdings. An important assumption is that asset holdings cannot fall below zero; no borrowing is permitted. Given this market structure, a household with positive wealth responds to a fall in household income by temporarily dis-saving. This means that households which have drawn a high proportion of low values for labor productivity in the recent past tend to have lower asset holdings in equilibrium than households which have typically enjoyed high productivity. The no-borrowing constraint is important because it limits the ability of low-wealth households to smooth consumption in the face of falls in their disposable income.

The second respect in which the economy differs from the simplest growth model is that there is a government which finances constant government spending by issuing one period debt and levying taxes. Contrary to the assumption in

[4] Weicher 1997 investigates the position of households with negative net worth in some detail. He finds that these households tend to have higher incomes and more assets than other poor households. In 1992 only 11.8 percent of those households with negative net worth (or 0.57 percent of the total population) had net worth of less than -$10,000.

Aiyagari and McGrattan (1998), the tax level is stochastic, meaning that households face risk at the aggregate as well as at the idiosyncratic level. The presence of aggregate risk means that in equilibrium there is inter-temporal variation in the shape of the joint distribution over productivity and wealth. From the households' perspective, debt and capital are perfect substitutes since the one period return to both is risk free, and there are no transaction costs. An equilibrium condition is that aggregate asset holdings at each date must equal the sum of the capital stock and the stock of outstanding government debt.

Individual states

A household's effective labor supply depends both on the hours it works and on its household-specific labor productivity, which is stochastic. At any date t, a household's productivity takes one of l values in the set E. Each household's productivity evolves independently according to a first-order Markov chain with transition probabilities defined by the $l \times l$ matrix Π. The probability distribution at t over E is represented by a row vector $p_t \in \mathrm{R}^l$, where $p_t \geq 0$ and $\sum_{i=1}^{l} p_{it} = 1$. If the probability distribution at date 0 is given by p_0 the distribution at t is given by $p_t = p_0 \Pi^t$. Given certain assumptions (which will be satisfied here) E has a unique ergodic set with no cyclically moving subsets and $\{p_t\}_{t=0}^{\infty}$ converges to a unique limit p^* for any p_0. Thus, given a population of measure 1, we can reinterpret p_t as the mass of the population in each productivity state at date t. I assume that $p_0 = p^*$, and impose an appropriate normalization such that $\sum_{i=1}^{l} p_i^* e_i = 1$.

There are two assets in this economy (capital and government debt) but by assumption they will pay the same return state-by-state. Thus the household effectively has a single savings instrument. Let A be the set of possible values for a household's holdings of this asset. I assume that a household's wealth at the start of period 0, denoted a_{-1}, is non-negative and that households are never able to borrow. Thus $A \subset \mathrm{R}_+$. Let $(A, \mathcal{A})$ and $(E, \mathcal{E})$ be measurable spaces where $\mathcal{A}$ denotes the Borel sets that are subsets of A and $\mathcal{E}$ is the set of all subsets of E. Let $e^t = \{e_0, ..., e_t\}$ denote a partial sequence of productivity shocks from date 0 up to date t, and let $e_t(e^t)$ denote the last element of this sequence. Let $(E^t, \mathcal{E}^t)$, $t = 0, 1, ...$ denote product spaces, and define probability measures

$$\mu^t : \mathcal{E}^t \rightarrow [0, 1], \; t = 0, 1, ... \tag{3.1}$$

where, for example, $\mu^t(e^t)$ is the probability of individual history e^t.

Aggregate states

The aggregate state of the economy at date zero, z_0, is defined by a measure $\lambda : \mathcal{A} \times \mathcal{E} \rightarrow [0, 1]$ describing the distribution of households across individual wealth and individual productivity at time 0, and the date 0 level of government debt B_{-1}.[5]

The only source of aggregate uncertainty in the model is the stochastic process for the economy-wide tax rate. This means that (given z_0) the aggregate state of the economy at t can be described by the history of the tax rate from date 0 up to and including date t. I call this object the aggregate history to date t, and denote it h^t. Let $\tau_t(h^t)$ denote the last element of this sequence. Let $\left(h^t, \mathcal{H}^t\right)$, $t = 0, 1, ...$ denote product spaces, and define probability measures

$$\nu^t : \mathcal{H}^t \rightarrow [0, 1]\,,\ \ t = 0, 1, ... \tag{3.2}$$

where, for example, $\nu^t(h^t : z_0)$ is the probability of aggregate history h^t. I shall use the notation $h^t \succeq h^{t-1}$ to indicate that h^t is a possible continuation of h^{t-1}.

The household's problem

The timing convention is that household productivity and the tax shock are observed before decisions are made in period t. In period 0, given the individual and aggregate states (a_{-1} and z_0) and the initial realizations for productivity and the tax rate ($e_0 = e_0(e^0)$ and $\tau_0 = \tau_0(h^0)$), the household chooses labor supply, savings and consumption for each possible sequence of individual productivity shocks and aggregate tax shocks. Let the sequences of measurable functions

$$\left.\begin{array}{l} n_t : E^t \times H^t \rightarrow [0, 1] \\ a_t : E^t \times H^t \rightarrow A \\ c_t : E^t \times H^t \rightarrow \mathrm{R}_+ \end{array}\right\} \ t = 0, 1, ... \tag{3.3}$$

describe this plan, where, for example, $a_t(h^t, e^t : a_{-1}, z_0)$ denotes the choice for savings that will be implemented at t if the aggregate history to date t is h^t and the individual history is e^t. Note that choices for consumption and labor supply have to be non-negative after every history, and labor supply cannot exceed the total time endowment which is equal to 1.

Expected discounted lifetime utility is given by

$$\sum_{t=0}^{\infty} \beta^t \sum_{h^t \in H^t} \nu^t(h^t) \sum_{e^t \in E^t} \mu^t(e^t) u\left(c_t\left(h^t, e^t\right), n_t\left(h^t, e^t\right)\right) \tag{3.4}$$

where β is the subjective discount factor. For the benchmark version of the model, I assume that the period utility function has the form introduced by Greenwood,

[5] The dependence of aggregate variables on z_0 and the dependence of household specific variables on a_{-1} are henceforth generally suppressed in the interests of brevity.

Hercowitz and Huffman (1988):

$$u(c,n)=\frac{1}{1-\gamma}\left[\left(c-\psi\frac{n^{1+1/\varepsilon}}{1+1/\varepsilon}\right)^{1-\gamma}-1\right]. \tag{3.5}$$

Here γ is the coefficient of relative risk aversion and ε is the inter-temporal (Frisch) elasticity of labor supply.[6]

The pre-tax real return to supplying one unit of effective labor at date t is given by the measurable function $w_t : H^t \to \mathrm{R}$. Similarly, the net one-period pre-tax return to one unit of the asset purchased at $t-1$ after history h^t is $r_t(h^t)$. The tax rate at t is assumed to take one of two possible values, $\tau_t(h^t) \in T = \{\tau_l, \tau_h\}$. In the benchmark version of the model, taxes are proportional, and apply equally to both asset and labor income. Thus the household budget constraints are given by

$$\begin{aligned} c_t\left(h^t,e^t\right)+a_t\left(h^t,e^t\right) &= \left[1+\left(1-\tau_t(h^t)\right)r_t(h^t)\right]a_{t-1}(h^{t-1},e^{t-1})+ \\ &\quad \left(1-\tau_t(h^t)\right)w_t(h^t)e_t(e^t)n_t\left(h^t,e^t\right) \end{aligned} \tag{3.6}$$

for all $e^t \in E^t$ such that $e^t \succeq e^{t-1}$, for all $h^t \in H^t$ such that $h^t \succeq h^{t-1}$, for $t = 0, 1, ...$, and where $a_{-1}(h^{-1}, e^{-1}) = a_{-1}$.

The solution to the household's problem is a set of decision functions (3.3) that maximize 3.4 taking as given (i) the household budget constraints (3.6), (ii) the price and tax functions w_t, r_t and τ_t, (iii) the probability measures (3.2 and 3.1), and (iv) the initial state (a_{-1}, z_0).

Labor supply

The utility function given in 3.5 has the convenient property that the labor supply choice is independent of the consumption / savings choice. In particular, assuming an interior solution, optimal individual labor supply is a simple function of the household-specific after-tax real return to working:

$$n_t(h^t,e^t:z_0,a_0)=\left[\frac{w_t(h^t:z_0)e_t(e^t)(1-\tau_t(h^t))}{\psi}\right]^{\varepsilon}.$$

Note that optimal labor supply does not depend on a_0, or on the history of productivity shocks up to $t-1$. Note also that the choice for ε determines the responsiveness of labor supply to variations in the household-specific real wage. An additional reason to use this functional form in the context of a model with

[6] The utility function is only defined for $c \geq 0$, $n \geq 0$, and $c \geq \psi\frac{n^{1+1/\varepsilon}}{1+1/\varepsilon}$.

heterogenous agents is that $N_t(h^t)$, equilibrium aggregate effective labor supply following history h^t, is a simple function of the inherited aggregate capital stock $K_{t-1}(h^{t-1})$, the current economy wide tax rate $\tau_t(h^t)$, the set of productivity shocks E, and the time-invariant distribution across these shocks p^* (see eq. 6.11 in the appendix for the derivation).

Production

Aggregate output after history h_t, $Y_t(h^t)$, is produced by competitive firms according to a Cobb-Douglas technology:

$$Y_t(h^t) = K_{t-1}(h^{t-1})^{\alpha} N_t(h^t)^{1-\alpha} \qquad h^t \succeq h^{t-1}$$

where $K_{t-1}(h^{t-1})$ denotes the capital stock in place at the start of period t, and $\alpha \in (0,1)$. Output can be transformed into private consumption, government consumption, and new capital according to

$$C_t(h^t) + G_t(h^t) + K_t(h^t) = Y_t(h^t) + (1-\delta)K_{t-1}(h^{t-1}) \qquad h^t \succeq h^{t-1}$$

where $C_t(h^t)$ denotes aggregate private consumption, $G_t(h^t)$ denotes government consumption, and $\delta \in [0,1]$ is the rate of depreciation.

Government

Real government spending is assumed constant and equal to G. Real government debt issued at date t is denoted $B_t(h^t)$. For any history h^t, this debt is assumed to pay a pre-tax one period real return equal to the economy-wide rate of return $r_t(h^t)$. Moreover, income from debt and capital are taxed at the same rate, implying that households are indifferent between saving in the form of capital or debt. Let aggregate asset holdings at the start of period $t+1$ be given by $A_t(h^t)$. The government's budget constraint is

$$\begin{aligned} & B_t(h^t) + \tau_t(h^t)\left[r_t(h^t)A_{t-1}(h^{t-1}) + w_t(h^t)N_t(h^t)\right] \\ = & \left(1 + r_t(h^t)\right) B_{t-1}(h^{t-1}) + G \qquad h^t \succeq h^{t-1} \end{aligned} \tag{3.7}$$

where $B_{-1}(h^{-1}) = B_{-1}$.

The process for taxes

The observation that the effects of current tax changes cannot be studied independently of the future tax changes that they imply is at the heart of the Ricardian equivalence proposition. However, even if government spending is held constant, many different paths for taxes are consistent with a stationary debt to GDP ratio.

The approach taken in this paper is to impose exogenous constant bounds on the level of debt issued by the government in the period, $B_t(h^t) \in D = [D_l, D_h]$, and to assume that the tax rate follows a Markov process such that if initial debt lies in the set D, then future debt always remains within D. This is implemented by ensuring that debt is always falling when $\tau = \tau_h$ and always rising when $\tau = \tau_l$, and by specifying transition probabilities such that for values of $B_t(h^t)$ close to D_h the probability of the high tax is always 1, while for $B_t(h^t)$ close to D_l it is always 0.[7] There is evidence that this is a reasonable specification for taxes. In particular, Bohn (1998) finds that the U.S. government has historically responded to increases in the debt-GDP ratio by raising the primary surplus, and that the debt-GDP ratio is mean-reverting once one controls for war-time spending and cyclical fluctuations.

Let $\pi_\tau : T \times D \times T \rightarrow [0,1]$ denote the time invariant transition probability function for taxes, where $\pi_\tau((\tau, B), \tau')$ is the probability that next period's tax rate is τ' given that the current tax rate is τ and the amount of new debt issued is B. The specification for π_τ adopted is as follows:

	$B \leq \underline{D}$	$\underline{D} < B < \overline{D}$	$B \geq \overline{D}$
$\pi_\tau((\tau_h, B), \tau_h)$	0	$\left[\frac{B-\underline{D}}{\overline{D}-\underline{D}}\right]^\lambda$	1
$\pi_\tau((\tau_l, B), \tau_l)$	1	$\left[\frac{\overline{D}-B}{\overline{D}-\underline{D}}\right]^\lambda$	0

where $\underline{D}$ and $\overline{D}$ are simple functions of D_h and D_l, and $\lambda \in (0, 1]$.

One feature of this specification is that the expected duration of a low tax regime is decreasing in the indebtedness of the government, while the expected duration of a high tax regime is increasing in B. The parameter λ controls the persistence of tax levels. If $\lambda = 1$, then the probability distribution over next period's tax rate is independent of the current rate. Reducing λ reduces the probability of a change in tax levels, conditional on a particular value for B.

The utility function (eq. 3.5) implies that aggregate labor supply is a increasing function of aggregate capital and a decreasing function of the tax rate. Thus a large capital stock improves the government's fiscal position via three channels: (i) more capital by itself implies more output and tax revenue, (ii) more capital raises the marginal product of labor, implying more labor supply and a further increase in output, and (iii) more capital implies a higher capital / labor ratio (see eq. 6.11), and thus lower interest payments on government debt. It is immediate that the government's fiscal position is also improved the lower is outstanding government debt, and the higher is the current tax rate (assuming

[7] Dotsey and Mao 1997 take a similar approach.

we are on the left side of the Laffer curve). Let $\kappa = [K_l, K_h]$ denote a set such that in equilibrium aggregate capital always lies in this set.[8] Taken together, the preceding observations imply that sufficient conditions for the upper bound on debt D_h not to be violated are:

$$\tau_h \geq \frac{r\left(K_l, N(K_l,\tau_h)\right) D_h + G}{r\left(K_l, N(K_l,\tau_h)\right)\left(D_h + K_l\right) + w\left(K_l, N(K_l,\tau_h)\right) N\left(K_l,\tau_h\right)} \tag{3.8}$$

and

$$\overline{D} \leq \frac{D_h - G + \tau_l\left[w\left(K_l, N(K_l,\tau_l)\right) N_l(K_l,\tau_l) + r\left(K_l, N(K_l,\tau_l)\right) K_l\right]}{1 + r\left(K_l, N(K_l,\tau_l)\right)\left(1-\tau_l\right)}. \tag{3.9}$$

where $N : \kappa \times T \to [0,1]$ is given in eq. 6.11.

The first condition says that conditional on the tax level being *high*, debt is non-increasing for all values for inherited debt $B \in D$ and for all values for inherited capital $K \in \kappa$. The second condition says that for all levels of inherited debt consistent with a *low* current tax level (i.e. $\forall B < \overline{D}$), new debt issued does not exceed D_h.

Similar conditions guarantee that the lower bound on debt D_l is not violated. The parameterization section describes how values are assigned to D_h, D_l, τ_h, τ_l and λ while ensuring that the conditions guaranteeing boundedness are satisfied.

Equilibrium

I define an equilibrium for this economy in the appendix.

3.1. Calibration

The model period is one year, the most appropriate horizon for considering tax changes. All parameter values are reported in annual terms in tables 1 and 2. The parameters relating to aggregate production are standard: capital's share in the production function α is set equal to 0.36 and the depreciation rate is 0.1. The risk aversion parameter in the utility function, γ, is set to 1, and the discount factor, β, is 0.96.

The intertemporal (Frisch) elasticity of labor supply parameter, ε, is an important parameter, and a somewhat controversial one (see Blundell and MaCurdy 1999 for a survey). MaCurdy (1981) estimates this elasticity to be in the range 0.1 to 0.45 for prime-age males. Blundell, Meghir and Neves (1993) study married women in the U.K. and estimate Frisch labor supply elasticities in the 0.5 to

[8] Appropriate values for K_l and K_h are determined within the numerical solution procedure.

1.0 range. I use a value of 0.3, which is lower than the value of 1.7 adopted by Greenwood, Hercowitz and Huffman (1998). Given the form of the utility function, labor supply is not affected by the level of non-labor income or the marginal utility of wealth. This means that in the model the uncompensated Marshallian wage elasticity is the same as the Frisch elasticity. Thus the fact that previous estimates of uncompensated elasticities are typically somewhat smaller than those for Frisch elasticities is one reason to pick a relatively low value for ε.[9] There is, moreover, little evidence of large labor supply responses to the changes in marginal tax rates that occurred during the 1980s (see Slemrod and Bakija 2000 for a discussion). Given the value for ε, the parameter ψ is set so that aggregate effective labor supply is equal to 0.3.

The household productivity process

The response of aggregate variables to tax changes is likely to depend on the distribution of wealth in the model economy, and in particular on the fraction of households on or close to the no-borrowing constraint. In the model described above, heterogeneity is generated endogenously as a consequence of households receiving uninsurable idiosyncratic productivity shocks. Thus the specification of the process for these shocks is critical.

I follow Domeij and Heathcote (2001) in searching for an income process with two broadly-defined properties. The first property is that the labor income uncertainty households experience is consistent with empirical estimates from panel data, so that the model is able to deliver appropriate time series variability in household income and consumption. The second is that the model economy generates realistic heterogeneity in terms of the distribution of wealth, and in particular, comes close to replicating the bottom tail of the observed wealth distribution.

I assume that l, the number of elements in the set E, is equal to three, since I find this to be the smallest number of states required to match overall U.S. wealth concentration and at the same time reproduce the fact that the wealth-poorest two quintiles hold a positive fraction of total wealth. Thus $E = \{e_1, e_2, e_3\}$, where the subscripts 1, 2 and 3 denote low, medium and high productivity respectively. I also assume that households cannot move between the high and low productivity levels directly, that the fraction of high productivity households equals the fraction of low productivity households, and that the probabilities of moving from the

[9] Note, however, that given a baseline value of 0.15 for the Frisch elasticity, MaCurdy estimates that hours worked are virtually unresponsive to changes in permanent non-wage income, virtually unresponsive to temporary income changes associated with temporary wage changes, and only mildly responsive to income changes associated with permanent wage changes. These findings broadly support using the Greenwood et. al. functional form for period utility.

medium productivity state into either of the others are the same. Thus the matrix Π is defined by just two parameters: $\Pi_{1,1}$ and $\Pi_{2,2}$, where $\Pi_{i,j}$ denotes the probability of transiting from state i to state j.

$$\Pi = \begin{bmatrix} \Pi_{1,1} & 1-\Pi_{1,1} & 0 \\ \frac{1-\Pi_{2,2}}{2} & \Pi_{2,2} & \frac{1-\Pi_{2,2}}{2} \\ 0 & 1-\Pi_{1,1} & \Pi_{1,1} \end{bmatrix} \tag{3.10}$$

Once mean productivity has been normalized to unity, the productivity process is therefore defined by a total of four free parameters: two levels and two transition probabilities.

Various authors have estimated stochastic AR(1) processes for logged household labor productivity and / or household income using data from the PSID. Such a process may be summarized by the serial correlation coefficient, ρ, and the standard deviation of the innovation term, σ. Allowing for the presence of measurement error and the effects of observable characteristics such as education and age indicates a ρ in the range 0.88 to 0.96, and a σ in the range 0.12 to 0.25.[10] I therefore impose two restrictions on the finite state Markov process for productivity: (i) that the first order autocorrelation coefficient equals 0.9, and (ii) that the variance for productivity is $0.05/(1-0.9^2)$, corresponding to a standard deviation for the innovation term in the continuous representation of 0.224. These are very close to the point estimates of Flodén and Lindé (1999), who consider a model with a labor supply choice and therefore focus explicitly on an exogenous process for labor productivity rather than labor income.

Ensuring that productivity shocks have the appropriate persistence and variance pins down two of the four productivity process parameters. I then adjust the remaining two free parameters to seek to match two properties of the empirical asset holding distribution: the Gini coefficient and the fraction of aggregate wealth held by the two poorest quintiles of the population. The second criterion is important because the households whose consumption is most sensitive to temporary tax changes are likely to be those with very low levels of wealth. Using data from the 1992 Survey of Consumer Finances, Diaz-Gimenez, Quadrini and Rios-Rull (1997) report a wealth Gini of 0.78, and find that the two poorest quintiles of the distribution combined hold 1.35 percent of total wealth.

The calibration procedure, described in Domeij and Heathcote (2001), delivers parameter values that satisfy all four criteria. This finding is interesting in light of the debate as to whether uninsurable fluctuations in earnings can account for

[10] See, for example, Card 1991, Hubbard, Skinner and Zeldes 1995 and Storesletten, Telmer and Yaron 1999. Heaton and Lucas 1996 allow for permanent but unobservable household-specific effects, and find a much lower ρ of 0.53, and a σ of 0.25.

U.S. households' wealth accumulation patterns.[11] Two key features of the process that reproduces the extreme wealth concentration observed in the United States are as follows. First, the fraction of households in the high productivity state at any point in time is small: 5.3 percent of the population. Second, the levels of productivity are asymmetric, in that e_3/e_2 is larger than e_2/e_1. Taken together, these two features imply that in equilibrium a small fraction of the population ends up holding a large fraction of total wealth, as is the case in the U.S.[12]

Table 3 provides a detailed comparison between the asset holding distribution observed in the data, and the average distribution observed over a long simulation of the calibrated benchmark model. The only respect in which the model does a relatively poor job is in terms of accounting for the substantial wealth holding of the richest 1 percent of households. Table 3 also reports the correlations between wealth, pre-tax labor earnings, and pre-tax income. The correlation between earnings and wealth is of particular interest, since it is those agents with both low wealth and low productivity who are most likely to be borrowing-constrained. This correlation is 0.36 in the model, versus 0.23 in the data. Figure 1 contains cumulative density functions describing the average (simulation) distribution of asset holdings across the entire population and conditional distributions given particular values for household productivity.

Adding up fixed private capital and the stock of durables owned by consumers, Aiyagari and McGrattan (1998) report a capital to annual output ratio of 2.5. Note that (by chance) the model reproduces this figure exactly. Given the choices for capital's share, the depreciation rate, and tax rates, this implies an average annual real after tax return to saving of 3.0 percent, a reasonable compromise for an economy in which stocks and bonds pay the same rate of return.

The tax process

All other model parameters relate to fiscal policy. The tax system in this model is represented by a single flat rate tax that applies equally to capital and labor income.[13] For agents who are not borrowing constrained, it is the marginal tax rate that is important for savings and labor supply decisions. However, for households for whom the constraint is binding, it is the average tax rate that

[11] See Quadrini and Rios Rull 1997 for a review of alternative theories of wealth inequality. Krusell and Smith 1998 find that their specification for idiosyncratic productivity shocks delivers a Gini co-efficient for wealth of only 0.25. They therefore intrroduce idiosyncratic shocks to the subjective discount factor as an additional mechanism for generating wealth inequality.

[12] On average, low, medium and high productivity types devote respectively 17, 27, and 43 percent of their time endowments to market work.

[13] In reality, the tax that a household pays is a complicated function of its income, and of the source of this income. See Altig and Carlstrom 1999 or Castañeda, Díaz-Giménez and Ríos-Rull 2000 for examples of treatments of non-linear tax schedules.

determines the level of consumption, given a choice for labor supply. Since the primary focus of the paper is on borrowing constraints as a propagation mechanism, I calibrate to average rather than marginal tax rates. Because there is a single tax rate in the model, the appropriate empirical average tax rate is the ratio of total government receipts to GDP.

The mean ratio of total (federal plus state and local) annual government current receipts to GDP in the United States between 1946 and 1999 was 0.26.[14] As is well known, this ratio has grown through time, from 0.23 in 1946 to 0.30 in 1999. Since there is no long-run growth in the size of government in the model, I first remove a linear trend from the revenue to GDP series in the data before computing the volatility and autocorrelation of the series. The detrended annual series has a standard deviation of 0.009 and autocorrelation equal to 0.63.[15] The average ratio for total government debt to GDP over the period 1946 to 1996 was 0.67.[16] In 1946 the value was 1.36; the post-war low of 0.47 was achieved in 1979.

There are six parameter values to be determined: the value for constant government consumption G, tax rates τ_l and τ_h, bounds on government debt D_h and D_l, and the persistence parameter λ. These parameter values are chosen simultaneously to approximately satisfy six criteria: (i) the average ratio of tax revenue to GDP in the model is 0.26, (ii) the first order autocorrelation of the ratio of tax revenue to GDP is 0.63, (iii) the standard deviation of the ratio of tax revenue to GDP is 0.009, (iv) the average ratio of government debt to GDP is 0.67, (v) high tax and low tax regimes are equally persistent, and the unconditional probability of being in either regime is 0.5, and (vi) debt remains bounded for every possible history for tax rates h^t.

The mean ratios for tax revenue and debt to GDP across a $10,000$ period simulation, along with the standard deviation and autocorrelation for tax revenue are reported in table 3.[17] In this simulation, the average duration of a tax change is 4.9 years.

[14] Data on tax revenue and GDP is from the National Income and Product Accounts, Tables 1.1 and 3.1, published by the Bureau of Economic Analysis.

[15] The Congressional Budget Office has estimated a series for the effective total federal tax rate. The mean and standard deviation of the 'all families' series between 1977 and 1999 are respectively 22.9 percent and 0.009.

[16] Data on debt is from the Statistical Abstract of the United States published by the Census Bureau. Data for 1996, for example, are from table no. 493 in the 2000 edition of the Abstract.

[17] Details of a numerical procedure that delivers parameter values with the desired properties are given in the computational appendix, available on request.

3.2. Alternative economies

Previous quantitative work suggests that weakening inter-generational links in an otherwise Ricardian world does not produce large real effects from temporary tax changes (see the introduction). This leaves capital market imperfections and distortionary taxation as the two primary candidates for generating big deviations from Ricardian equivalence. The benchmark model described above features both borrowing constraints and distortionary taxes. The first two alternative economies I consider treat these two features separately. In the first, taxes are lump-sum, isolating the market structure as a propagation mechanism. In the second, markets are complete, isolating the effects of distortionary taxation. The third economy considers a variation on the benchmark economy in which labor supply is perfectly inelastic, in order to gauge the sensitivity of results to the intertemporal elasticity of labor supply. The final economy constitutes an extension of the benchmark model designed to capture some features of life-cycle consumption and savings behavior.

The parameters listed in table 1 are held constant across the various alternative economies. Other parameters relating to the household productivity process and the process for taxes are recalibrated for each new economy following a procedure analogous to the one described above for the benchmark model. In particular, the household productivity parameters are adjusted so that each model economy reproduces the targeted features of the U.S. wealth distribution, and the tax parameters are adjusted so that tax revenue has the same persistence and variance as has been observed empirically. These parameter values are given in table 2. Properties of the implied asset holding distributions and tax processes are reported in table 3.

Lump sum tax economy

The first economy is identical to the benchmark model except that taxes are lump sum rather than proportional. Although lump-sum taxes are unrealistic in practice, this is an interesting model to consider because in a world of lump sum taxes and infinitely-lived households, the Ricardian Equivalence proposition would obtain were asset markets complete. Thus any real effects from temporary tax changes in this economy will be directly attributable to the presence of borrowing constraints coupled with uninsurable risk. The lump sum tax economy is therefore an attractive framework for assessing the potential importance of capital market imperfections as a propagation mechanism.

One small difference in the calibration approach relative to the benchmark economy pertains to the use of tax revenue. To ensure that low productivity, low wealth households can realize a positive marginal utility of consumption in the presence of lump-sum taxes, I assume that the government makes constant lump-sum transfers ϕ to households, and that government consumption is always zero.

The household budget constraint in this case is therefore given by

$$c_t\left(h^t,e^t\right)+a_t\left(h^t,e^t\right)=\left(1+r_t(h^t)\right)a_{t-1}(h^{t-1},e^{t-1})+w_t(h^t)e_t(e^t)n_t\left(h^t,e^t\right)+\phi-\tau_t(h^t)$$

Government debt now evolves according to

$$B_t(h^t)+\tau_t(h^t)=\left(1+r_t(h^t)\right)B_{t-1}(h^{t-1})+\phi. \tag{3.11}$$

Complete markets economy

In the complete markets economy, markets exist which allow households to fully insure against idiosyncratic productivity risk, and against any distributional effects from aggregate tax shocks. I therefore adopt the representative agent abstraction, and assume that the representative household's labor productivity is constant and equal to unity. As in the benchmark economy, taxes are proportional to income. Since the representative agent will always choose positive asset holdings, the only source for Ricardian non-neutrality in this economy arises from the fact that taxes are distortionary.

Exogenous labor supply economy

This economy is identical to the benchmark economy, except that labor supply is exogenous. I assume that each household supplies 0.3 units of labor per period.

Life cycle economy

This economy is designed to capture the idea that some households may be borrowing constrained simply because they are young and at the bottom of an upward-sloping lifetime earnings profile. I model the life-cycle in a highly stylized fashion, which allows me to consider the life-cycle economy as a special case of the benchmark model in which the transition probability matrix for household productivity shocks is suitably modified. In particular, I assume that on top of the labor productivity risk described above, households face aging risk. For a household with low productivity, aging amounts to transiting to the medium productivity state. A medium productivity household who ages transits to the high productivity state. A high productivity household who ages transits to the low productivity state. This last event may be thought of as an elderly agent dying and being replaced by a newborn successor who inherits all the financial assets of the parent, but none of the parent's human capital. Transition probabilities for this economy, described by the matrix $\widehat{\Pi}$ are constructed as follows.

First, I assume that the probability of transiting from state e_i via the mechanism identified as aging is equal to $1/(\widehat{p}_i L)$, where $\widehat{p}_i$ is the fraction of the population with productivity e_i in the ergodic distribution over E, and L is a constant. Note that a fraction $\widehat{p}_3/(\widehat{p}_3 L)$ of the population dies and is replaced in each period, implying that L may be interpreted as expected lifetime. The events of aging and receiving a productivity shock are assumed mutually exclusive. The overall probability of moving from state i to state j is therefore equal to the probability of transiting from i to j via aging, plus the probability of transiting from i to j via a productivity shock, conditional on not aging.

To calibrate the life cycle model I generate $\widehat{\Pi}$ using exactly the same Π matrix as in the benchmark model:[18]

$$\widehat{\Pi} = \begin{pmatrix} 0 & \frac{1}{\widehat{p}_1 L} & 0 \\ 0 & 0 & \frac{1}{\widehat{p}_2 L} \\ \frac{1}{\widehat{p}_3 L} & 0 & 0 \end{pmatrix} + \begin{pmatrix} (1-1/\widehat{p}_1 L) & 0 & 0 \\ 0 & (1-1/\widehat{p}_2 L) & 0 \\ 0 & 0 & (1-1/\widehat{p}_3 L) \end{pmatrix} \Pi.$$

Note that the fractions $\widehat{p}_i$ are the solutions to the system of equations $\widehat{p} = \widehat{p}\widehat{\Pi}$.

While the aging / productivity shock distinction is a convenient conceptual device, it is irrelevant for agents in the model who only care about the implied transition probability matrix $\widehat{\Pi}$. Note that relative to the benchmark economy, low and medium productivity agents now attach higher probability to a productivity increase. Thus we may expect these households to exhibit lower demand for precautionary savings, and for their consumption to be more tax-sensitive. This is the sense in which the life-cycle economy is designed to emphasize the importance of the no-borrowing constraint.

3.3. Numerical solution

It is known to be difficult to solve for an equilibrium in economies with heterogeneous agents, incomplete markets, and aggregate uncertainty. I therefore adopt the strategy proposed by Krusell and Smith (1998).[19] In particular, I

[18] The estimates for the variance and persistence of labor income risk used to calibrate the benchmark model are based on household level data that has been purged of variation attributable to age and education. Thus these estimates should be compared to the properties of the process for productivity implied by the Π matrix (rather than the $\widehat{\Pi}$ matrix). I construct $\widehat{\Pi}$ using the original Π matrix from the benchmark model because I was unable to find an alternative data-consistent specification for Π that generates realistic wealth inequality when households take as given the process associated with $\widehat{\Pi}$.

[19] Den Haan 1997 proposes a similar algorithm. Other papers to implement the Krusell and Smith approach include Storesletten et. al. 1998 and Castaneda et. al. 1998.

assume that when solving their problems, rather than using all of the information about the aggregate state of the economy contained in h^t, households instead only consider the information contained in $Z_t = \left(K_{t-1}(h^{t-1}), B_{t-1}(h^{t-1}), \tau_t(h^t)\right)$. Given this assumption, I consider a recursive formulation of the household's problem in which households take as given a law of motion for aggregate capital $G : \kappa \times D \times T \to \kappa$. The solution to the household's problem is a decision rule of the form $a' : E \times A \times \kappa \times D \times T \to A$. Given decision rules, the economy is simulated forward, and a regression is run on the simulated data to update the co-efficients in the forecasting rule G. This procedure is repeated until convergence, at which point the forecasting rule G that households take as given is such that their behavior generates a law of motion for capital for which the best predictor function (of the same functional form as G) is precisely the forecasting rule G.[20]

Figure 2 contains the benchmark economy equilibrium decision rules for consumption and net savings, given each possible combination of household-specific productivity and the economy-wide tax rate.[21] Consumption is an increasing function of wealth, while net savings is decreasing in wealth. Low productivity households are universally dis-savers, while high productivity households are net savers except at very high levels of wealth. This is the natural pattern given that households wish to smooth consumption in the presence of mean-reverting shocks to income. Medium productivity households are net savers at very low levels of wealth, and dis-savers at higher levels. For households with high productivity, the optimal consumption and savings rules are close to linear in wealth, while for less productive types, the marginal propensity to consume out of wealth is decreasing in wealth. That this is attributable to the presence of the no-borrowing constraint is evidenced by the fact that non-linearities are most pronounced at very low levels of wealth, and for households with the lowest value for productivity. For example, given zero wealth and low productivity, the marginal propensity to consume out of wealth is one.

It is important to note that the only reason why these decision rule may not represent the solution to the original household problem is that households may forecast future aggregate capital with error under the forecasting rule G. Thus, in evaluating the cost of using a limited information set (and a linear forecasting

[20] The revised household problem, the numerical procedure for solving this problem given a specification for the forecasting rule G, and the numerical implementation of the Krusell and Smith procedure for iterating on the parameters of the forecasting rule to minimuze forecasting errors are all described in a computational appendix which is available on request.

[21] In figure 2 aggregate capital and debt are set to their average equilibrium levels. Mean household wealth in equilibrium is the sum of aggregate capital and aggregate debt. To magnify non-linearities, decision rules are plotted only for low to moderate values for household wealth.

rule) it is important to examine the magnitude of these forecasting errors and the implied errors in forecasting future factor prices. For the models considered in this paper, the differences between actual future prices and forecasted future prices are very small and on the order of those encountered by Krusell and Smith. For example, the cumulative forecasting error for the net pre-tax interest rate (the marginal product of capital minus the depreciation rate) at a ten year horizon rarely amounts to more than two tenths of a percent difference between the predicted value for the interest rate (conditional on the realized sequence for taxes) and the actual value observed in the simulation. As an alternative metric, the absolute difference between the predicted rate and the realized rate over a ten year horizon is rarely more than one basis point (one hundredth of a percentage point).[22] Thus it is hard to imagine that improving forecasting accuracy would lead to large changes in individual decision rules or the aggregate behavior of the economy.

Why is it that higher moments of the wealth distribution do not seem to be very useful for forecasting future prices? The intuition is similar to that given in Krusell and Smith. Firstly, note that if the saving rule were exactly linear in wealth then redistributing wealth among agents with a particular productivity realization would have no effect on aggregate savings. In light of the shape of the savings rules in figure 2, redistributing wealth between agents with moderate or greater wealth will therefore have little effect on aggregate savings. Recall, however, that at low levels of wealth the marginal propensity to save out of wealth is increasing in wealth. This suggests that wealth redistributions between very poor households and richer households could significantly change total savings. There are several reasons why this is not an important problem in practice. First, the shape of the wealth distribution does not change much through time; for example, from figure 3 the poorest 40 percent of households almost invariably account for between 0.8 percent and 1.8 percent of aggregate wealth. Second, households with low savings propensities account for a disproportionately small fraction of aggregate economic activity in general and aggregate consumption in particular. Third, other variables in the forecasting rule for capital contain information that partially substitutes for more detailed information about the shape of the wealth distribution. For example, the correlation in a simulation between the Gini coefficient for asset holdings and the level of debt is -0.90.

4. Results

Benchmark economy - aggregate effects of tax changes

[22] More details on forecasting accuracy are in the computational appendix.

Consider the aggregate effects of tax changes across a $10,000$ period simulation of the benchmark economy during which values for the tax rate are drawn according to the specified stochastic process.[23] The focus of the paper is on the response of aggregate consumption to tax changes. One statistic which is informative in this regard is the change in consumption between two consecutive periods characterized by different tax rates, relative to the change in tax revenue between the same two periods. A value of zero for this statistic, which I call the propensity to consume out of income tax (PCT), would indicate that in aggregate households behave in a Ricardian fashion, and adjust private saving rather than consumption in response to tax changes. Figure 3 contains scatter plots of the PCT for the 1019 periods during the simulation in which the tax rate went up and the 1018 periods in which the tax rate fell. On average, aggregate consumption fell by 28.6 cents for every dollar increase in tax revenue following tax increases, and fell by 28.9 cents for every dollar of revenue lost following tax cuts (see table 4). These responses are much larger than those typically found in models that generate real effects by shortening the household's horizon (see the introduction). There is little difference between the average response of aggregate consumption to tax decreases versus tax increases. There is, however, some variation through time: the largest response of consumption to a tax cut in the simulation is 31.6 cents per dollar change in tax revenue while the smallest response is 25.7 cents (see figure 3).

To assess the persistence of consumption responses, I report the average responses at horizons of one and four years, in addition to the immediate impact effect (see table 4). In the year of a tax increase, aggregate consumption was on average 20.2 cents below its long-run mean value for every dollar of taxes collected in excess of mean revenue. At dates one year after tax increases occurred, consumption was on average 13.1 cents below its mean value for every such surplus tax dollar collected in the previous year. Four years after a tax increase was observed, consumption was close to its long-run average value.

To isolate the various propagation channels for tax shocks and thereby assess their relative quantitative importance I compare the benchmark economy to the various alternative economies described above. For each economy I first compute an average joint distribution over wealth and productivity across a long simulation. I then compute the impact response to a tax cut (increases are symmetric) starting with this distribution and with the average simulation stock of

[23] An initial joint distribution across individual states was taken from an economy without aggregate uncertainty (see the available computational appendix). The full-blown economy was then simulated for $11,000$ periods before computing statistics for the last $10,000$ periods of the sample.

government debt (see table 5). As a reference point, such a tax cut in the benchmark model increases aggregate consumption by 28.7 cents for every dollar of tax revenue lost.

For each alternative economy, I also examine behavior at the household level, since households with different asset holding / productivity characteristics may exhibit very different responses to tax changes. The figures reported in table 6 are ratios of the difference in household consumption across the two tax rates relative to the difference in household tax payments. These differences are computed given the average joint distribution over wealth and productivity, and given the mean debt to GDP ratio. Consider, for example, low productivity households with zero wealth in the benchmark model. For these households, the difference in optimal consumption across the two tax regimes is $1.28 for each dollar difference in taxes paid.

Lump-sum tax economy

Consider first the economy with lump-sum taxes. In this economy, neither individual nor aggregate labor supply is directly affected by changes in the level of a lump-sum tax. Moreover the capital stock is fixed in the short run. Thus a tax change will have no immediate effect on output or factor prices, and for any given household, after-tax income will change by the same dollar amount as the tax level. If households were to respond to changes in public saving by making exactly off-setting adjustments to private saving then debt and tax finance would be equivalent.

Table 6 indicates that the consumption of households with both very low wealth and low productivity varies one for one with the lump-sum tax level. Such households do no saving irrespective of the tax rate. Low and medium productivity households with less than median wealth are sufficiently concerned with maintaining a buffer stock of asset holdings to self insure against income shocks that they do not completely offset tax changes by adjusting private saving. As wealth increases, however, the gap between optimal consumption in the two tax regimes narrows, indicating that households are increasingly willing to use their assets to consumption-smooth through tax shocks.

High productivity households actually consume more when the tax level is high than when it is low, conditional on a given level of asset holdings. To understand this, recall that the interest rate is endogenous. The flip side of the increase in aggregate consumption following a tax cut is a fall in investment. This translates into higher expected future interest rates. High productivity or high wealth households (who assign a low probability to the possibility of being borrowing-constrained in the near future) respond to the increase in the expected return to saving by increasing saving and reducing consumption.

The aggregate propensity to consume out of a tax change is on the order of 12 cents per dollar change in tax revenue. This figure is one measure of the potency of liquidity constraints as a propagation mechanism for tax shocks. It would be larger if the interest rate were exogenous, in which case all households would increase consumption in response to a tax cut (see Heathcote 1999).

Complete markets economy

In the economy with complete markets, households face no risk of becoming borrowing constrained. Thus all real effects from tax shocks derive from the fact that taxes are distortionary. These distortionary effects are large. Table 5 indicates that given mean values for capital and debt, a switch from the high tax rate to the low rate is associated with a 23 cent increase in consumption for every dollar of tax revenue lost. In the complete markets economy this tax cut leads to an immediate 0.97 percent increase in aggregate labor supply, and a 0.60 percent increase in output. The increase in output is absorbed by a 0.86 percent rise in investment, and a 0.74 percent rise in private consumption (recall that government consumption is assumed fixed). Comparing the complete markets economy to the benchmark incomplete markets model, the responses of aggregate labor supply and output to a tax change are very similar across the two economies. However, while investment responds more strongly than consumption in the complete markets model, the reverse is true in the benchmark model. This accounts for the finding that the propensity to consume out of income tax in the complete markets model is 5.7 cents less per dollar than for a similar tax cut in the benchmark economy.

A low income tax rate in the current period makes leisure expensive relative to consumption, and thus increases labor supply (see eq. 6.11). Given the utility function, the marginal utility of consumption rises which tends to increase consumption for the representative household. At the same time, since tax rates are persistent, a low current tax rate signals a high expected return to saving, and thus provides an incentive to increase saving and reduce current consumption. The fact that consumption increases following a tax cut indicates that the increase in the after-tax return to working is quantitatively the more important effect.

Within the complete markets model it is easy to compare alternative specifications for the utility function. Suppose, for example, that the utility function has the form $u(c,n) = \left[c^{\mu}(1-n)^{1-\mu}\right]^{1-\sigma}/(1-\sigma)$ where $1/\sigma$ is the intertemporal elasticity of substitution. Given a value for μ such that on average $n = 0.3$, and recalibrating parameters defining the tax process such that tax shocks have the observed variance and persistence, I compute the impact effect of a typical tax cut for two alternative values for the σ. For $\sigma = 1$, the PCT and percentage changes in labor supply, consumption and investment are respectively 20.6 cents per dollar, 5.0 percent, 0.6 percent and 12.5 percent. For $\sigma = 3$, the corresponding figures are 47.3 cents per dollar, 3.9 percent, 1.5 percent, and 7.0 percent. The message from this exercise is that the baseline specification for preferences (eq. 3.5) implies moderate real effects from tax shocks relative to alternative plausible specifications.

Exogenous labor supply economy

In this economy tax changes have effects both via the borrowing constraint channel, and also because changes in the tax rate affect the after-tax return to saving. As in the calibrated benchmark economy, the stationary distribution over productivity shocks is such that the vast majority of households are in the medium productivity state. With exogenous labor supply, optimal consumption of a medium productivity / median wealth household in the high tax state is only 7 cents lower for each extra dollar of taxes paid (see table 6). This compares to a 17 cent difference in the lump-sum tax economy, and accounts for the fact that, contrary to the lump-sum tax model, aggregate consumption barely moves in response to a tax shock (see table 5).

The reason for the relative insensitivity of consumption to the tax rate is twofold. First, taxes are proportional in this economy, so that compared to the lump-sum case, tax shocks imply much smaller changes in after-tax income for low income households. Thus it is easier for these households to consumption-smooth through tax shocks. The second reason is that a tax increase in this economy reduces the expected after-tax return to saving, because tax shocks are persistent. Thus all households have an incentive to increase consumption following a tax increase, which works in the opposite direction of the borrowing constraint effect.

Note that high productivity households consume less in the high tax state than in the low tax one. I interpret this as reflecting a wealth effect. A high productivity household prefers a low tax rate when income is high (the current period) and higher tax rates in the future when expected labor productivity (and thus taxable income) is lower.

Life-cycle economy

The life-cycle economy was designed to generate large consumption responses to tax changes by modelling the interaction of a borrowing constraint with an upward-sloping earnings profile. I find, however, that the responses of all aggregate variables to tax shocks are very similar to those in the benchmark model. One might suspect that this reflects the fact that this model generates too little wealth inequality - the Gini co-efficient for wealth is only 0.63 compared to a value of 0.78 for the other economies. The reason for the low Gini, however, is that the rich are too poor, and not that the poor are too rich; 24 percent of households in this economy have less than 2 percent of mean wealth.

Upward income mobility means that for any given level of wealth, low and medium productivity households assign a lower probability to being borrowing constrained in the near future. As a device for magnifying the importance of the no-borrowing constraint, however, this is a double-edged sword. Expecting high future labor income, households on or very close to the no-borrowing constraint are more willing to consume out of tax cuts. At the same time, households with higher levels of wealth are less concerned about the possibility of the constraint binding in the future, and are more willing to smooth through tax shocks by adjusting their asset holdings.

Benchmark economy: Cross-sectional variation in responses to tax shocks

Given the discussion above, it is now straightforward to interpret the decision rules of different types of household to tax changes in the benchmark economy (see table 6 and figure 2).

For high productivity households, the borrowing constraint is of little concern, since these households are net savers, and assign low probability to the event of receiving two bad shocks in the near future and thereby transiting to the low productivity state. The primary channel through which tax shocks affect the behavior of these households is by altering their incentive to work, and thereby their marginal utility of consumption. Thus these households work and consume more in the low tax state than in the high tax state, irrespective of their level of asset holdings.

As in the lump-sum tax and exogenous labor economies, consumption is most tax sensitive among households with low or medium productivity and low levels of asset holdings. For households with low productivity and zero wealth, the no-borrowing constraint is binding for both tax rates. The difference in consumption across tax regimes is larger than the difference in tax payments because labor supply and pre-tax income is lower when taxes are high. At higher levels of wealth, however, the low productivity household's decision rules corresponding to the high and low tax rates lie virtually on top of each other. This is because these households work less than more productive households and consume primarily out of untaxed wealth rather than taxed income. Thus they are more affected by changes in the incentive to save (which tend to reduce consumption in the low tax state) rather than changes in the incentive to work.

The decision rules for medium productivity households constitute an intermediate case relative to the high and low productivity types considered thus far. Consumption is 'excessively' sensitive to the tax rate at low levels of wealth, but the borrowing constraint is never binding for these households, and thus even at zero wealth, medium productivity households respond to tax changes by adjusting both consumption and savings. At higher levels of wealth, consumption becomes less tax-sensitive, as households become increasingly willing to adjust asset holdings in order to consumption smooth in the face of tax shocks.[24]

Benchmark economy: Longitudinal variation in responses to tax shocks

The top left panel of figure 3 shows that across a $10,000$ period simulation, tax decreases tend to have larger effects on aggregate consumption the smaller is the fraction of total wealth accounted for by the poorest 40 percent of households. During a prolonged interval of high tax rates the ratio of debt to GDP gradually falls. At the same time, households gradually run down their asset holdings and more households pile up against the no-borrowing constraint. This accounts for the positive relation between the wealth of poor households and the debt to GDP ratio (see the bottom left panel of figure 3). When the tax rate does eventually fall there is a large increase in aggregate consumption, both because poor households have largely exhausted their assets and also because given low debt the tax cut is expected to be relatively persistent.

Consider next a sustained period of low taxes. Expecting a tax increase in the future, households gradually accumulate more precautionary asset holdings. The longer taxes remain low, the larger is the ratio of debt to GDP when the tax rate eventually rises, and the larger the asset holdings of the wealth-poorest households. There is, however, no clear relation between the fraction of wealth held by these households and the size of the fall in aggregate consumption when taxes rise. On the one hand, if taxes have been low for a long time, fewer households are close to the borrowing constraint and unable to smooth through a tax increase. On the other hand, the ratio of debt to GDP will tend to be high, implying that the tax increase is relatively persistent and should therefore have a large effect on aggregate consumption.

4.1. Quantitative importance of distortions versus liquidity effects

Comparison of the various economies considered above provides a clear picture of what factors are important for generating large real effects from tax shocks. Distortionary taxation can imply large responses for aggregate consumption and

[24] There is a literature focussing on insurance effects which operate when missing insurance markets mean that distortionary tax changes affect the inter-temporal distribution of idiosyncratic risk (see Barsky, Mankiw and Zeldes 1986, Kimball and Mankiw 1989, Chan 1983, Fremling and Lott 1994, and Croushore 1996). In the model developed here these insurance effects are not quantitatively important since taxes are strictly proportional to income, and there are no transfers.

labor supply in response to a change in the proportional income tax rate. Of the two channels through which distortionary taxation operates, the effect of taxing labor income on the labor supply choice is likely quantitatively more important than the effect of taxing capital income on the savings choice. Comparing the benchmark model with the exogenous labor economy, the large difference in the response of aggregate consumption to a tax shock is striking given that the two models differ only with respect to the intertemporal labor supply elasticity. One reason endogenous labor supply is important when agents are heterogenous is that it effectively magnifies the importance of the borrowing constraint. In particular, a tax increase reduces the after-tax return to working and discourages labor effort. After-tax income therefore falls both because the tax rate increases and because pre-tax income falls. This means that low income low wealth households are thrown that much harder against the no-borrowing constraint, and are forced to reduce consumption.

The combination of missing insurance markets coupled with borrowing constraints implies large responses at the household level for households with very low wealth and low or medium productivity. However, these do not translate into particularly large effects at the aggregate level. For example, the aggregate propensity to consume out of a tax cut in the benchmark economy is only six cents larger than in the complete markets version. Given the extreme concentration of wealth in the model, why is it that the borrowing constraint has only moderate effects at the aggregate level?

Several reasons have already been discussed. In particular, relative to a specification with lump-sum taxes, proportional tax shocks imply relatively small and easily-smoothed changes in disposable income for low income / low wealth households. Moreover, the facts that capital income is taxed proportionately and that tax shocks are persistent works against a large aggregate consumption response, since a tax decrease increases the expected after-tax return to saving. Two additional factors that mitigate against the quantitative importance of the no-borrowing constraint are as follows.

First, the households whose consumption is most tax sensitive are also the households with the lowest levels of equilibrium consumption. For example, households with medium productivity and no wealth consume 58 cents less for each extra dollar of taxes paid in the high tax regime but enjoy a consumption level of only around 60 percent of mean per-capita consumption. While the bottom quintile of households ranked by consumption account on average for only 10.8 percent of aggregate consumption (compared to 7.5 percent in the U.S. in 1988; Cutler and Katz 1992) the top quintile accounts for 44.8 percent of the total (compared to 37.2 percent in the U.S.). Thus the decisions of wealthier households for whom the borrowing constraint is of little concern are disproportionately important at the aggregate level.

Second, idiosyncratic productivity shocks are more persistent than aggregate tax shocks in the calibrated model economy. This suggests it is relatively easy to consumption smooth through tax shocks by appropriately adjusting asset holdings through time, and relatively hard to consumption smooth through productivity shocks. In Heathcote (1999), a working paper version of this paper, I find that increasing the persistence of tax shocks does in fact increase the responsiveness of aggregate consumption to lump-sum tax shocks. Assessing the effects of less persistent income shocks is more difficult, since reducing persistence tends to simultaneously reduce wealth inequality and thereby weakens liquidity effects. In Heathcote (1999) I consider a specification in which the annual autocorrelation co-efficient for productivity shocks is 0.6 compared to the value of 0.9 adopted here, and generate wealth inequality by introducing persistent idiosyncratic shocks to the household discount factor. In that economy, the real effects of tax shocks are larger than those reported here; for example with exogenous labor supply and proportional taxes, the propensity to consume out of income tax is 8.3 cents per dollar, as compared to 1.9 cents in the corresponding economy here. However, caution is appropriate in attributing this difference to the different persistence of productivity shocks, since it is difficult to assess the role of taste shocks in generating large real effects.

5. Conclusions

In the model economies studied here, income tax changes have real effects both because they distort labor supply and savings decisions, and also because missing insurance markets coupled with a borrowing constraint limit households' ability to smooth consumption through time. The response of aggregate consumption to simulated tax shocks is typically large, and consistent in sign and magnitude with many empirical estimates of the effects of historical tax changes in the United States. A temporary tax increase in the benchmark model economy reduces aggregate consumption by around 29 cents for every additional dollar of tax revenue raised. At the aggregate level the effects of tax changes are largest in models that incorporate both the distortionary effects of proportional taxes and the liquidity effects that result from an incomplete markets environment. Most of the response of aggregate consumption in these models, however, is attributable to the fact that flat rate taxes are distorting.

Aggregate measures of the effects of tax changes hide wide variation in the response at the household level. I find that the consumption of low income low wealth households is the most sensitive to changes in the tax level, and that this sensitivity is primarily attributable to the no-borrowing constraint. Among wealthier households, tax shocks affect behavior primarily because they change the expected time paths for the returns to working and saving. Except for households receiving little labor income, changes in the return to working have larger effects on household consumption than changes in the return to saving. There is therefore scope for future work to explore alternative ways of introducing labor supply in the household utility function.

In terms of the quantitative importance of the no-borrowing constraint, the nature of the tax change is key. If a proposed tax change implies a small absolute change in the incomes of households close to the constraint (as is the case for a change in a flat rate income tax) the liquidity constraint will have relatively small effects at the aggregate level. If a tax change implies a large absolute change in these households income (as under a change in a lump-sum tax or a change in some type of welfare benefit) the effect on aggregate consumption will be much larger.

What fraction of households are typically on or close to a borrowing constraint? The greater the fraction of such households, the larger will be the effects of tax changes. The approach to answering this question taken here is to assume that no borrowing is permitted and to calibrate the model to reproduce the United States wealth distribution. One could argue that in reality the constraint is effectively looser than the one imposed here (if, for example, a degree of non-collateralized borrowing is possible) or tighter (if, for example, certain types of wealth such as consumer durables are too illiquid to be readily adjusted to smooth through income shocks).

In a simulation of the benchmark model there are typically many wealth-poor households (as in the United States). The presence of the no-borrowing constraint therefore implies that averaging across the population, the mean percentage change in household consumption following a tax change is also large. At the same time, richer households account for a disproportionately large fraction of total consumption, and the consumption of these households is less sensitive to changes in the tax rate. This suggests a possible explanation for why empirical work based on micro data has often found larger effects from tax changes than are apparent in aggregate data.

References

[1] Aiyagari, S. Rao. 1994. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109, 659-684.

[2] Aiyagari, S. Rao and Ellen R. McGrattan. 1998. The optimum quantity of debt. *Journal of Monetary Economics*, 42, 447-469.

[3] Altig, David and Charles T. Carlstrom. 1999. Marginal tax rates and income inequality in a life-cycle model. *American Economic Review,* 89, no. 5, 1197-1215.

[4] Altig, David and Steve J. Davis. 1989. Government debt, redistributive fiscal policies, and the interaction between borrowing constraints and intergenerational altruism. *Journal of Monetary Economics*, 24, 3-29.

[5] Auerbach, Alan J. and Laurence J. Kotlikoff. 1987. *Dynamic Fiscal Policy*, Cambridge University Press.

[6] Barksy, Robert B., N. Gregory Mankiw and Stephen P. Zeldes. 1986. Ricardian consumers with Keynesian propensities. *American Economic Review*, 76, no. 4, 676-690.

[7] Barro, Robert J. 1974. Are government bonds net wealth? *Journal of Political Economy*, 82, no. 6, 1095-1117.

[8] Bernheim, B. Douglas. 1987. Ricardian equivalence: An evaluation of theory and evidence. *NBER Macroeconomics Annual*, 263-304.

[9] Blinder, Alan S. 1981. Temporary income taxes and consumer spending. *Journal of Political Economy*, 89, no. 1, 26-53.

[10] Blundell, Richard W. and Thomas E. MaCurdy. 1999. Labor supply: A review of alternative approaches. In O. Ashenfelter and D. Card eds., Handbook of Labor Economics, Vol 3A. Elsevier Science, Amsterdam.

[11] Blundell, Richard W., Costas Meghir, and P. Neves. 1993. Labour supply: An intertemporal substitution. *Journal of Econometrics*, 59, 137-160.

[12] Braun, Anton. 1994. Tax disturbances and real economic activity in the postwar United States. *Journal of Monetary Economics*, 33, 247-280.

[13] Cardia, Emanuela. 1997. Replicating Ricardian equivalence tests with simulated series. *American Economic Review*, 87, no. 1, 65-79.

[14] Castañeda, Ana, Javier Diaz-Gimenez, and Jose-Victor Rios-Rull. 1998. Exploring the income distribution business cycle dynamics. *Journal of Monetary Economics*, 42, 93-130.

[15] Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull. 2000. Accounting for earnings and wealth inequality. Mimeo, University of Pennsylvania.

[16] Chan, Louis Kuo Chi. 1983. Uncertainty and the neutrality of government financing policy. *Journal of Monetary Economics*, 11, 351-372.

[17] Christiano, Lawrence J. and Jonas D. M. Fisher. 1997. Algorithms for solving dynamic models with occasionally binding constraints. Working Paper 97-11. Federal Reserve Bank of Cleveland.

[18] Cox, Donald and Tullio Jappelli. 1993. The effect of borrowing constraints on consumer liabilities. *Journal of Money, Credit and Banking*, 25, no. 2, 197-213.

[19] Croushore, Dean. 1996. Ricardian equivalence with wage-rate uncertainty. *Journal of Money, Credit and Banking*, 28, no. 3, 279-293.

[20] Cutler, David M. and Lawrence F. Katz. 1992. Rising inequality? Changes in the distribution of income and consumption in the 1980's. *American Economic Review*, 82, no. 2, 546-551.

[21] Daniel, Betty C. 1993. Tax timing and liquidity constraints: A heterogeneous-agent model. *Journal of Money, Credit and Banking*, 25, no. 2, 176-196.

[22] Den Haan, Wouter. 1997. Solving dynamic models with aggregate shocks and heterogeneous agents. *Macroeconomic Dynamics*, 1, no. 2.

[23] Diaz-Gimenez, Javier, Vincenzo Quadrini, and Jose-Victor Rios-Rull. 1997. Dimensions of inequality: Facts on the U.S. distributions of earnings, income and wealth. *Quarterly Review (Spring)*. Federal Reserve Bank of Minneapolis.

[24] Dotsey and Mao. 1997. The effects of fiscal policy in a neoclassical growth model. Working Paper 97-8, Federal Reserve Bank of Richmond.

[25] Feldstein, Martin. 1988. The effects of fiscal policies when incomes are uncertain: A contradiction to Ricardian equivalence. *American Economic Review*, 78, no. 1, 14-23.

[26] Flodén, Martin and Jesper Lindé. 1999. Idiosyncratic risk in the U.S. and Sweden: Is there a role for government insurance?. *Review of Economic Dynamics*, forthcoming.

[27] Fremling, Gertrude M. and John R. Lott, Jr. 1994. Do deficits affect the level of insurance? *Journal of Money, Credit and Banking*, 26, no. 4, 934-940.

[28] Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman. 1988. Investment, capacity utilization, and the real business cycle. *American Economic Review*, 78, no. 3, 402-417.

[29] Hayashi, Fumio. 1985. The effect of liquidity constraints on consumption: A cross-sectional analysis. *Quarterly Journal of Economics*, February, 183-206.

[30] Heaton, John and Deborah Lucas. 1996. Evaluating the effect of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy,* 104, no. 3, 443-487.

[31] Hubbard, Glenn R. and Kenneth L. Judd. 1986. Liquidity constraints, fiscal policy, and consumption. *Brookings Papers on Economic Activity*, 1, 1-59.

[32] Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. Precautionary saving and social insurance. *Journal of Political Economy*, 103, no. 2, 360-399.

[33] Jappelli, Tulio. 1990. Who is credit constrained in the U.S. economy? *Quarterly Journal of Economics,* February, 219-234.

[34] Kimball, Miles S. and N. Gregory Mankiw. 1989. Precautionary saving and the timing of taxes. *Journal of Political Economy*, 97, no. 4, 863-879.

[35] Krusell, Per and Anthony A. Smith Jr. 1998. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106, no. 5, 867-896.

[36] MaCurdy, Thomas E. 1981. An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy*, 89, 1059-1085.

[37] McGrattan, Ellen. 1994. The macroeconomic effects of distortionary taxation. *Journal of Monetary Economics*, 33, 573-601.

[38] Modigliani, Franco and Charles Steindal. 1977. Is a tax rebate an effective tool for stabilization policy? *Brookings Papers on Economics Activity*, 1:1977, 175-209.

[39] Parker, Jonathan A. 1999. The reaction of household consumption to predictable changes in social security taxes. *American Economic Review*, 89(4), 959-973.

[40] Poterba, James M. 1988. Are consumers forward looking? Evidence from fiscal experiments. *American Economic Review*, 78, no. 2, 413-418.

[41] Poterba, James M. and Lawrence H. Summers. 1987. Finite lifetimes and the effects of budget deficits on national saving. *Journal of Monetary Economics*, 20, 369-391.

[42] Seater, John J. 1993. Ricardian equivalence. *Journal of Economic Literature*, vol. 31, 142-190.

[43] Shapiro, Matthew D. and Joel Slemrod. 1995. Consumer response to the timing of income: Evidence from a change in tax withholding. *American Economic Review*, 85, no. 1, 274-283.

[44] Slemrod, Joel and Jon Bakija. 2000. Taxing ourselves: A citizen's guide to the great debate over tax reform. Second Edition, The MIT Press.

[45] Storesletten, Kjetil, Chris Telmer and Amir Yaron. 1999. Asset pricing with idiosyncratic risk and overlapping generations. Mimeo, Carnegie Mellon University.

[46] Souleles, Nicholas S. 1999. The response of household consumption to income tax refunds. *American Economic Review*, 89, no. 4, 947-958.

[47] Souleles, Nicholas S. 2001. Consumer response to the Reagan tax cuts. *Journal of Public Economics*, forthcoming.

[48] Trostel, Philip A. 1993. The nonequivalence between deficits and distortionary taxation. *Journal of Monetary Economics,* 31, 207-227.

[49] Wilcox, David W. 1989. Social security benefits, consumption expenditure, and the life cycle hypothesis. *Journal of Political Economy*, 97, no. 2, 288-304.

[50] Weicher, John. 1997. The rich and the poor: Demographics of the U.S. wealth distribution. *Federal Reserve Bank of St. Louis Review*, 79, no. 4, 25-37.

[51] Woodford, Michael. 1990. Public debt as private liquidity. *American Economic Review*, 80, no. 2, 382-388.

[52] Zeldes, Stephen P. 1989. Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy*, 97, no. 2, 305-346.

6. Appendix: definition of equilibrium

An equilibrium for the benchmark economy is a set of functions e_t, a_t, c_t, n_t, w_t, r_t, τ_t, K_t, B_t, N_t, C_t, Y_t, probability measures μ^t and ν^t, and an initial state $z_0 = (\lambda, B_{-1})$ such that $\forall h^t \in H^t$, $\forall e^t \in E^t$, $\forall a_{-1} \in A$ and $\forall t = 0, 1, ...$

1. a_t, c_t, and n_t solve the household maximization problem, described in the text.
2. $\{\mu^t(\cdot)\}_{t=0}^{\infty}$ is consistent with the transition probability matrix Π, i.e. for any $P = P_0 \times ... \times P_t \in \mathcal{E}^t$
$$\mu^t(P) = \sum_{i \text{ s.t. } e_i \in P_{t-1}} \mu^{t-1}(P_0 \times ... \times P_{t-2} \times e_i) \sum_{j \text{ s.t. } e_j \in P_t} \Pi_{ij} \tag{6.1}$$
where
$$\mu^0(P_0) = 1 \text{ if } e_0 \in P_0 \text{ and } 0 \text{ otherwise}$$
3. $\{\nu^t(\cdot : z_0)\}_{t=0}^{\infty}$ is consistent with the transition function for taxes π_τ, i.e. for any $P = P_0 \times ... \times P_t \in \mathcal{H}^t$.
$$\nu^t(P : z_0) = \sum_{h_t \in P} \nu^t(h^t : z_0)$$
where $\forall i \in [1, t]$, $\forall h^i \in P_0 \times ... P_i$
$$\nu^i(h^i : z_0) = \nu^i(h^{i-1} : z_0)\pi_\tau\left(\left(\tau_{i-1}(h^{i-1}), B_{i-1}(h^{i-1})\right), \tau_i(h^i)\right) \tag{6.2}$$
$$h^i \succeq h^{i-1} \tag{6.3}$$
and
$$\nu^0(h^0 : z_0) = 1 \text{ if } \tau_0 = h^0 \text{ and } 0 \text{ otherwise.}$$
4. Aggregate quantities are consistent with individual decision rules:
$$A_t(h^t : z_0) = \int_{A\times E} \sum_{e^t \in E^t} \mu^t(e^t) a_t(h^t, e^t : z_0, a_{-1}) d\lambda \tag{6.4}$$
$$C_t(h^t : z_0) = \int_{A\times E} \sum_{e^t \in E^t} \mu^t(e^t) c_t(h^t, e^t : z_0, a_{-1}) d\lambda \tag{6.5}$$
$$N_t(h^t : z_0) = \int_{A\times E} \sum_{e^t \in E^t} \mu^t(e^t) e_t(e^t) n_t(h^t, e^t : z_0, a_{-1}) d\lambda \tag{6.6}$$

5. The market for savings clears:

$$K_{t-1}(h^{t-1}) + B_{t-1}(h^{t-1}) = A_{t-1}(h^{t-1}) \tag{6.7}$$

where $B_{-1}(h^{-1}) = B_{-1}$ and $A_{-1} = \int_{A\times E} a_{-1} d\lambda$

6. Factor markets clear:

$$r_t(h^t) = \alpha \left[K_{t-1}(h^{t-1})\right]^{\alpha-1} \left[N_t(h^t)\right]^{1-\alpha} - \delta \tag{6.8}$$

$$w_t(h^t) = (1-\alpha) \left[K_{t-1}(h^{t-1})\right]^{\alpha} \left[N_t(h^t)\right]^{-\alpha} \tag{6.9}$$

where $h^t \succeq h^{t-1}$ and $K_{-1}(h^{-1}) = A_{-1} - B_{-1}$.

7. The government budget constraint is satisfied and debt remains bounded:

$$B_t(h^t)+\tau_t(h^t) \left[r_t(h^t)A_{t-1}(h^{t-1}) + w_t(h^t)N_t(h^t)\right] = \left(1 + r_t(h^t)\right) B_{t-1}(h^{t-1})+G$$

$$B_t(h^t) \in [0, \infty) \tag{6.10}$$

where $h^t \succeq h^{t-1}$ and $B_{-1}(h^{-1}) = B_{-1}$.

8. The goods market clears.

$$C_t(h^t) + G + K_t(h^t) = Y_t(h^t) + (1-\delta)K_{t-1}(h^{t-1})$$

where $h^t \succeq h^{t-1}$ and $K_{-1}(h^{-1}) = A_{-1} - B_{-1}$.

Note that combining 6.6 and 6.9 we get

$$\begin{aligned} N_t(h^t \quad &: \quad z_0) = \int_{A\times E} \sum_{e^t\in E^t} \mu^t(e^t)e_t(e^t)n_t(h^t, e^t : z_0, a_{-1})d\lambda \qquad (6.11) \\ &= \int_{A\times E} \sum_{e^t\in E^t} \mu^t(e^t)e_t(e^t) \left[\frac{w_t(h^t)e_t(e^t)(1-\tau_t(h^t))}{\psi}\right]^{\varepsilon} d\lambda \\ &= \sum_{i=1}^{l} p_i^* e_i \left[\frac{(1-\alpha)\left[K_{t-1}(h^{t-1})\right]^{\alpha}\left[N_t(h^t)\right]^{-\alpha} e_i(1-\tau_t(h^t))}{\psi}\right]^{\varepsilon} \\ &= \left(\sum_{i=1}^{l} p_i^* e_i^{1+\varepsilon} \left[\frac{(1-\alpha)K_{t-1}(h^{t-1})^{\alpha}(1-\tau_t(h^t))}{\psi}\right]^{\varepsilon}\right)^{\frac{1}{1+\alpha\varepsilon}} \end{aligned}$$

Table 1. Parameter values common across economies

Preferences	β	0.96	*Production*	α	0.36
	γ	1.0		δ	0.10
	ε	0.3			

Table 2. Parameter values that vary across model economies

		ECONOMY				
		Bench	Lump-sum taxes	Complete markets	Exog. labor	Life cycle
Preferences	ψ	50.0	100.0	35.0	0.0	65.0
	e_1	0.1875	0.2389		0.1731	0.1803
	e_2	0.8483	0.8373		0.8472	0.8156
Productivity	e_3	4.3703	3.8369		4.6816	4.2018
	$\Pi_{3,3}$	0.9000	0.9001		0.9000	0.6762
	$\Pi_{2,2}$	0.9881	0.9843		0.9893	0.9691[i]
	G/Y	0.228	0.00[ii]	0.233	0.228	0.228
	τ_h	0.3432	τ_h/Y=0.2707[iii]	0.3325	0.3417	0.3441
Fiscal Policy	τ_l	0.3175	τ_l/Y=0.2491	0.3083	0.3176	0.3153
	D_h/Y	0.84	0.99	0.84	0.84	0.84
	D_l/Y	0.50	0.35	0.50	0.50	0.50
	λ	0.33	0.33	0.39	0.33	0.35

Table 3. Properties of average asset holding distribution over a 10,000 period simulation, mean simulation ratios of capital and debt, and simulation properties of ratio of tax revenue to GDP.[iv]

	DATA	*MODELS*				
		Bench.	Lump-sum taxes	Complete markets	Exog. labor	Life cycle
				Wealth Distribution		
Gini	0.78	0.78	0.78		0.78	0.63
99-100[v]	29.6	11.5	11.1		11.6	6.8
90-100	66.1	60.1	60.1		60.3	40.4
80-100	79.5	83.8	83.9		83.9	62.8
40-100	98.6	98.7	98.6		98.6	98.1
				Correlations		
earnings, income[vi]	0.93	0.96	0.98		0.91	0.98
earnings, wealth	0.23	0.36	0.41		0.32	0.14
income, wealth	0.32	0.60	0.58		0.68	0.36
				Ratios to GDP		
capital	2.50	2.50	2.87	2.23	2.42	2.42
debt	0.67	0.67	0.67	0.67	0.67	0.67
				Tax revenue to GDP		
mean	26.0	25.7	25.8	26.4	26.1	26.1
std dev	0.009	0.010	0.011	0.010	0.010	0.012
autocorrelation	0.63	0.64	0.61	0.62	0.63	0.61

Table 4. Responses to tax changes in benchmark economy: Averages over 10,000 period simulation, and impact effect given average joint distribution over wealth and productivity

	Means across simulation		*Average dist.*
	Tax Increases	Tax Decreases	Tax Decrease
	Pre-shock values		
debt to GDP	0.679	0.662	0.680
capital to GDP	2.492	2.508	2.507
tax revenue to GDP	0.248	0.267	0.267
wealth Gini	0.778	0.781	0.780
% < 2%[vii]	4.30	4.90	4.56
wealth poorest 40%[viii]	1.38	1.25	1.27
	Propensity to consume out of tax (PCT)		
$100 \times (C_t - C_{t-1}) / (T_t - T_{t-1})$	-28.6	-28.9	-28.7
	Deviation from mean consumption		
$100 \times (C_t - \mu_c) / (T_t - \mu_T)$	-20.2	-21.4	-26.4
$100 \times (C_{t+1} - \mu_c) / (T_t - \mu_T)$	-13.1	-15.1	
$100 \times (C_{t+4} - \mu_c) / (T_t - \mu_T)$	-3.6	-5.5	
	Percentage changes on impact		
tax revenue	7.14	-6.66	-6.66
GDP	-0.66	0.66	0.66
labor supply	-1.03	1.04	1.04
investment	-0.61	0.60	0.59
consumption	-0.97	0.99	0.98

Table 5. Responses to a tax cut in various economies (tax increase similar)[ix]

	ECONOMY				
	Bench.	Lump-sum taxes	Complete markets	Exog. labor	Life cycle
	Pre-shock values				
wealth Gini	0.780	0.780		0.780	0.633
% < 2%	4.56	5.98		4.54	24.3
wealth poorest 40%	1.27	1.23		1.32	1.91
	Propensity to consume out of tax (PCT)				
$100 \times (C_t - C_{t-1}) / (T_t - T_{t-1})$	-28.7	-12.2	-23.0	-1.9	-29.1
	Percentage changes on impact				
tax revenue	-6.66	-7.97	-6.59	-7.11	-7.48
GDP	0.66	0.02	0.60	0.00	0.74
labor supply	1.04	0.01	0.97	0.00	1.17
investment	0.59	-0.82	0.86	-0.13	0.61
consumption	0.98	0.37	0.74	0.07	1.12

Table 6. Consumption sensitivity to the tax rate for households with different asset holding / productivity characteristics (difference in consumption by type divided by difference in taxes paid by type)

	$[C_i(\tau = \tau_h) - C_i(\tau = \tau_l)] / [T_i(\tau = \tau_h) - T_i(\tau = \tau_l)]$		
		productivity (e_i)	
wealth (a_i)	low	medium	high
		Benchmark economy	
zero	-1.28	-0.58	-0.36
median (0.05 x mean)	-0.52	-0.40	-0.36
0.5 x mean	-0.10	-0.26	-0.35
mean	-0.02	-0.21	-0.34
2.0 x mean	0.02	-0.17	-0.33
5.0 x mean	0.00	-0.13	-0.31
		Lump sum tax economy	
zero	-1.00	-0.39	0.07
median (0.06 x mean)	-0.36	-0.17	0.08
mean	-0.02	0.03	0.11
		Complete markets economy	
representative agent		-0.25	
		Exogenous labor economy	
zero	-1.00	-0.23	-0.05
median (0.05 x mean)	-0.31	-0.07	-0.05
mean	0.10	0.05	-0.04
		Life-cycle economy	
zero	-1.28	-0.77	-0.36
median (0.49 x mean)	0.05	-0.25	-0.34
mean	0.13	-0.20	-0.33

---

[i] In the life cycle economy $\Pi_{1,1} = 0.6762$, $\Pi_{1,2} = 0.3238$, $\Pi_{1,3} = 0.0000$, $\Pi_{2,1} = 0.0058$, $\Pi_{2,2} = 0.9691$, $\Pi_{2,3} = 0.0251$, $\Pi_{3,1} = 0.2487$ $\Pi_{3,2} = 0.0751$, $\Pi_{3,3} = 0.6762$.

[ii] In the lump-sum tax economy I introduce constant (date and individual state invariant) lump sum transfers equal to 24 percent of average output. This is convenient for computational purposes, since it effectively guarantees that the marginal utility of consumption is positive, even for low productivity / low wealth households.

[iii] Ratio of lump-sum tax levels to GDP.

[iv] The empirical wealth distribution statistics are from Diaz-Gimenez et. al. 1997 and is based on data from the 1992 Survey of Consumer Finance. The statistics characterizing the empirical tax process are based on linearly detrended data from 1946 to 1999 (see text for details).

[v] Fraction of total economy assets held by wealth-richest 1 percent of households.

[vi] Income and earnings are pre taxes and transfers in the data and in the models.

[vii] Percentage of households with less than two percent of mean assets.

[viii] Percentage of total economy assets held by wealth-poorest 40 percent of households.

[ix] All the figures in tables 5 and 6 are computed given an average (over a 10,000 period simulation) joint distribution over asset holdings and productivity, and an average quantity of government debt.

Figure 1: Distribution over asset holdings
(average over 10,000 period simulation)

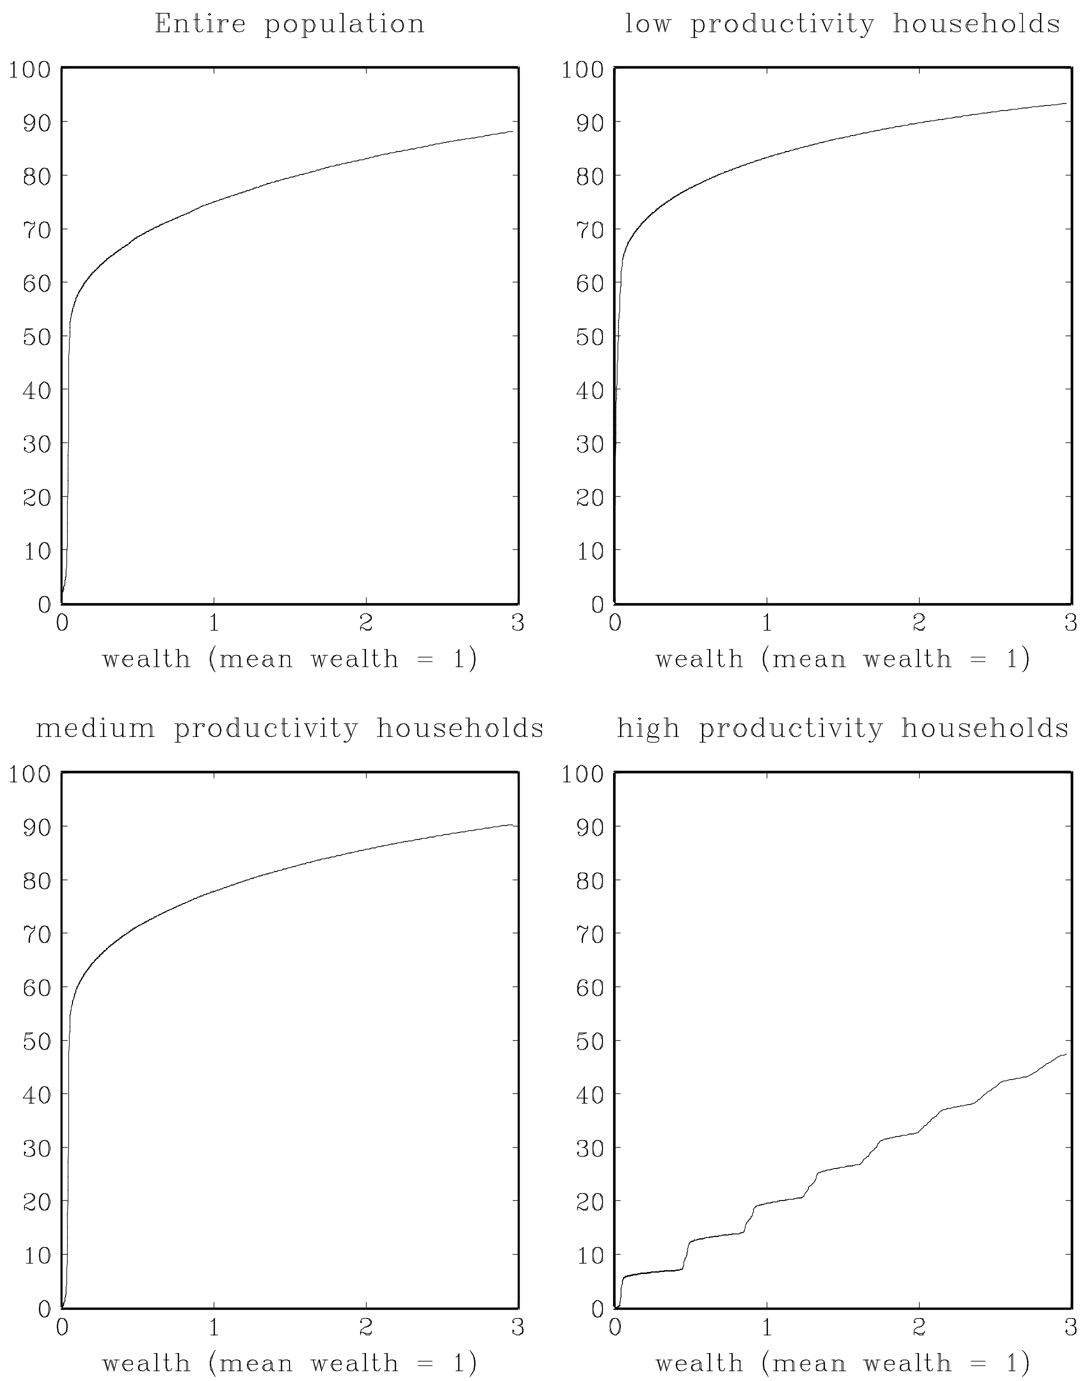

Figure 2: Decision rules for consumption and savings given mean aggregate debt and capital

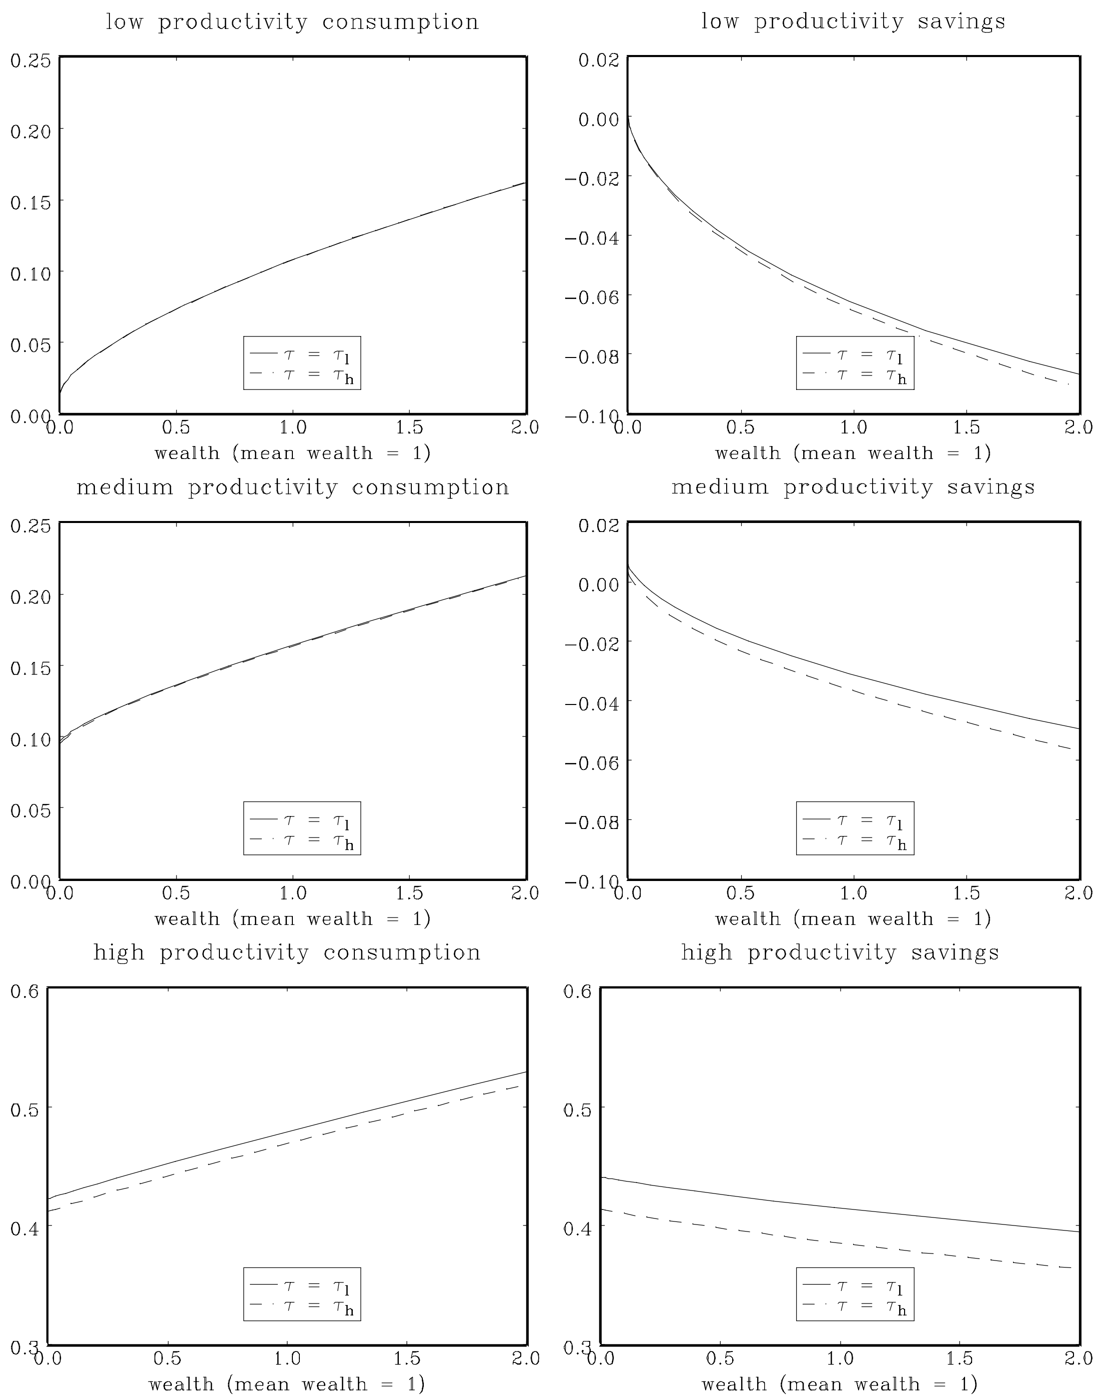

Figure 3: Distribution of changes in agg. consumption per dollar change in agg. tax revenue (PCT)