Government and Public Policy

Marina Azzimonti, Jonathan Heathcote, Kjetil Storesletten

1 Introduction

The government has a large impact on economic outcomes through fiscal policy, monetary policy, and regulation policy. In this chapter, we focus on fiscal policy; in particular, on taxes, government spending, and debt. After presenting a summary of how governments tax, spend and borrow in practice, we turn to theory and discuss how fiscal policy choices impact the competitive equilibrium allocation, and how to frame the problem of optimizing over those policy choices. Monetary policy is discussed in Chapters 15 and 16. Our discussion centers on developed economies, with a particular focus on the United States. Chapter 22 introduces fiscal policy in emerging markets.

In Chapter 6, we discussed conditions under which the First Welfare Theorem holds. When markets are complete and competitive and there are no public goods or externalities –i.e., there are no "market failures"– competitive equilibrium allocations are Pareto optimal. Why then, do governments intervene in the economy? There are three main rationales.

The first is that there are public goods, such as national defense, that the market cannot provide because there is no way to restrict the enjoyment of public goods to those households who choose to pay for them. There are other goods and services, like education and healthcare, that are not pure public goods, but whose consumption confers large positive externalities. For example, if my neighbors are vaccinated, they are less likely to make me sick. Thus, absent government involvement, education and healthcare might be under-consumed.

A second reason governments intervene is that markets are not complete and competitive, in part because of private information frictions. For example, it might be difficult to buy private unemployment insurance, or annuities that insure against longevity risk. Thus, there may be a role for the government to provide public unemployment insurance, or to fund a public pension system. It is also possible that absent government intervention, the economy might occasionally get stuck in an inefficiently depressed equilibrium because of frictions in private markets. Thus, the government intervened during the Global Financial Crisis in 2008, bailing out a range of financial institutions to avoid a cascade of bankruptcies, and cutting taxes to try to boost consumer confidence.

The third rationale for government intervention is redistribution. Market economies tend to generate substantial income inequality, as discussed in Chapter 9 (and later in Chapter 19). This inequality may be Pareto efficient, but taxing the rich in order to fund transfers to the poor will generate a more equal allocation of resources, and one that a majority of households might prefer. In many economies, transfers account for most of total government spending. We discuss redistribution in Section 7.1.

The impact of the government on equilibrium allocations depends not just on how much the government wants to spend, but also on how the government pays for that spending. In practice, the taxes that households and firms are required to pay depend on the choices they make about how much to work and earn, how much to consume versus save, and how much to invest. Thus the tax system distorts all those choices, ultimately reducing output. We explore the effects of distortionary taxation by adding proportional capital and labor income taxes to a standard neoclassical growth model. Note that higher public consumption or higher transfers necessitates higher tax rates, implying larger distortions to private sector choices and lower efficiency. There is significant cross-country variation in total government spending and in the extent of redistribution through the tax and transfer system. That suggests societies differ in how they view the trade-off between the benefits of a more equitable distribution of resources or higher public good provision, versus the efficiency costs of higher and more distortionary taxes (see Figure 1).

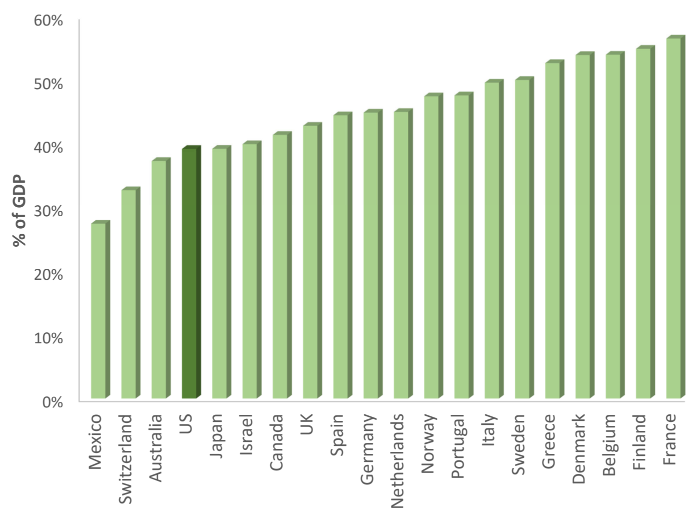

Figure 1: Government Spending across Countries (avg 2010-2019)

2 Public Finance: An Overview of the Data

The government spends money on publicly-provided goods and services G_t, such as education and defense, and makes transfers T_t to individuals and corporations, such as food stamps and agricultural subsidies. These expenditures are financed out of tax revenues Rev_t, collected through taxes on goods and services (sales and excise taxes), taxes on income (income and payroll taxes), property taxes, and taxes on corporate profits. When revenues are insufficient to cover expenditures, the government borrows from domestic households and firms or from international lenders. Denoting the stock of debt at the start of period t by B_{t-1}, net borrowing is equal to $B_t - B_{t-1}$. We denote the nominal interest rate on public debt by i_t, so $i_t B_{t-1}$ is interest payments. The government budget constraint can be written as

$$\underbrace{G_t + T_t + i_t B_{t-1}}_{\text{Expenditures}} = Rev_t + \underbrace{B_t - B_{t-1}}_{\text{Borrowing}}. \tag{1}$$

When expenditures – including interest payments – exceed revenues, we say that the government runs a *deficit*. In that case, $B_t > B_{t-1}$ and public debt rises. When expenditures are lower than revenues, the government runs a *surplus* and debt decreases. The stock of debt at any point in time, then, is the cumulative sum of net deficits run by a government through history. We now review some facts about the evolution of the main components of the government budget constraint to frame the topics discussed in this chapter.

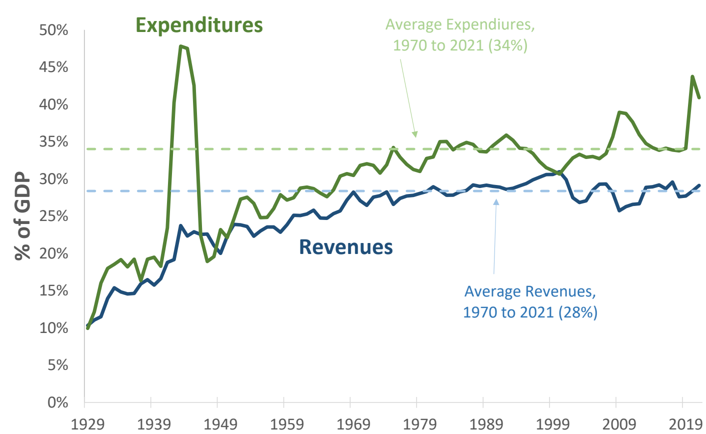

Figure 2: Revenues and Outlays, as percentages of GDP

Figure 2 shows revenues and expenditures as percentages of GDP for the U.S. between 1929 and 2021. The series, which incorporate all levels of government (federal, state, and local), were obtained from the NIPA tables constructed by the Bureau of Economic Analysis (see Appendix 8.1 for details). Three key points can be drawn from this figure. First, both revenue and spending exhibit upward trends between 1929 and 1970, and then stabilize. Second, expenditures tend to exceed revenues, implying that the U.S. government typically runs deficits. Between 1970 and 2021 expenditures and revenues averaged 34 percent and 28 percent of GDP respectively. Third, expenditures jump during periods of war or recession, while revenues typically fall. Large increases in expenditure are evident during World War II, the Great Recession of 2007-2009, and the COVID-19 recession in 2020.

Figure 3 describes how the sources of tax revenue in the United States have changed over time. Over the post-war period, income and social insurance tax revenues increased significantly, and now account for around 11 and 7 percent of GDP, respectively. Revenue from sales and import taxes have been relatively constant at about 8 percent of GDP, while revenue from corporate taxes has declined and is now less than 2 percent of GDP.

The theoretical literature on the impact of taxes differentiates between taxes on labor income, on capital income, and on consumption. In Section 3 we discuss the impact of these taxes on labor supply, investment, and savings choices. Note, however, that this simple theoretical categorization does not map cleanly into the empirical partition of taxes: in particular, while labor earnings are part of the base for U.S. personal income taxes, income taxes taxes also apply to income accruing to capital, including unincorporated business income, dividend, interest, and rental income, and capital gains.

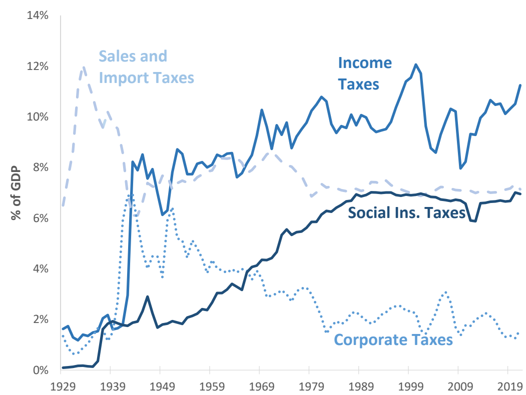

Figure 3: Taxes by Category, as percentages of GDP

The government can postpone taxation by using public debt to finance expenditure. Does it matter whether the government finances spending out of current taxation versus whether it issues debt which must be repaid out of future tax revenue? In Section 4 we show that when taxes are lump-sum, the timing of taxes is irrelevant. This famous result is known as "Ricardian Equivalence." However, in the more realistic case in which taxes are distortionary, the timing of taxes does matter. What is then the optimal timing of taxes? In Section 5, we formalize the problem of a benevolent government that chooses a sequence for taxes and for debt to maximize social welfare, following a formulation known as the "Ramsey problem." Solving this problem we demonstrate an important "tax smoothing" result: it is optimal to finance temporary shocks such as wars, recessions or pandemics mostly by issuing debt. Evidence that governments do in fact smooth taxes over time is presented in Figure 4 (left panel), showing the total deficit (expenditures minus revenues, solid line), the primary deficit (which is the deficit excluding interest payments, dark bars), and interest payments (light bars). The U.S. government borrowed heavily during wars and recessions, particularly during WWII and the COVID-19 pandemic. The right panel shows the stock of debt. In the U.S., most public debt is issued by the Federal government, the result of balanced budget rules written into State constitutions. In other countries, a significant part of borrowing is done by sub-national units.

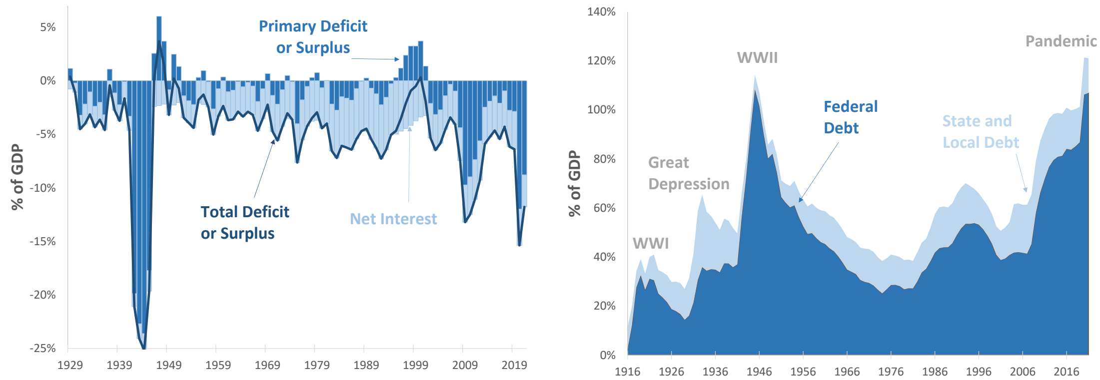

Figure 4: Left: Total Deficits, Primary Deficits, and Net Interest Outlays. Right: Total Debt. All as percentages of GDP.

While borrowing is largest during recessions and wars, we also see that the U.S. has run persistent deficits. This raises the question of how much debt is sustainable. One way to approach this question is to compare projections of future government deficits to the deficit levels that are consistent with a stable debt to GDP ratio.

Let D_t denote the primary deficit at date t (government spending excluding interest payments minus revenue). The government budget constraint (eq. 4) can then be written, in nominal terms, as

$$B_t = B_{t-1} \cdot (1 + i_t) + D_t.$$

Dividing through by nominal GDP at t gives

$$b_t = b_{t-1} \cdot \frac{1 + i_t}{(1 + \gamma_t)(1 + \pi_t)} + d_t,$$

where lower case letters denote values relative to nominal GDP, and where γ_t and π_t denote the growth rates of real GDP and the price level between $t-1$ and t. Let $1 + r_t = (1 + i_t)/(1 + \pi_t)$ denote the *ex post* gross real interest rate between $t-1$ and t. Thus, the debt to GDP ratio evolves according to

$$b_t = b_{t-1} \cdot \frac{1 + r_t}{1 + \gamma_t} + d_t.$$

It is clear from this equation that the value of the real interest rate relative to the real growth rate is critical for the dynamics of public finances. When the primary deficit d_t is zero, the debt to GDP ratio will rise when $r_t > \gamma_t$, and will fall when $r_t < \gamma_t$.

One can ask what size primary deficit is consistent with a constant debt to GDP ratio. Debt will rise over time ($b_t > b_{t-1}$) if and only if

$$d_t > \frac{\gamma_t - r_t}{1 + \gamma_t} \cdot b_{t-1}. \tag{2}$$

At the time of writing (September, 2023) U.S. government debt held by the public is around 100 percent of annual U.S. GDP – i.e., $b_{t-1} = 1.0$. The growth rate of real GDP in the United States varies over time, but has averaged around 3 percent per year in the post-War period, suggesting $\gamma_t = 0.03$. The interest rate on 10 year inflation-protected government bonds is currently a little over 2 percent, suggesting $r_t = 0.02$.

Plugging these numbers into our debt substainability equation suggests that the largest primary deficit consistent with debt not rising is approximately 1 percent of GDP. How does this compare to the actual primary deficit? The primary federal deficit in fiscal year 2022 was 3.6 percent of GDP, and the Congressional Budget Office (CBO) is forecasting primary deficits over the next 10 years of around 3.0 percent of GDP.[1] Thus, the U.S. debt to GDP ratio is likely to continue to grow. But will debt explode? Perhaps surprisingly, the arithmetic suggests not. In particular, given constant values for r and $\gamma > r$, any size primary deficit is consistent with a stable debt to GDP ratio, as long as that ratio is large enough. For example, suppose r_t and γ_t are expected to remain constant at values of 2 and 3 percent respectively. A 3 percent of GDP primary deficit is then consistent with a stable debt to GDP ratio of 309 percent of GDP (in terms of equation (2), $0.03 = \frac{0.01}{1+0.03} \times 3.09$). However, we should be very cautious about this calculation. As debt rises, the equilibrium real interest rate is likely to rise – investors will demand higher returns to buy all that debt. And once the differential between γ and r changes sign, stabilizing the debt to GDP ratio will require primary surpluses rather than deficits.[2]

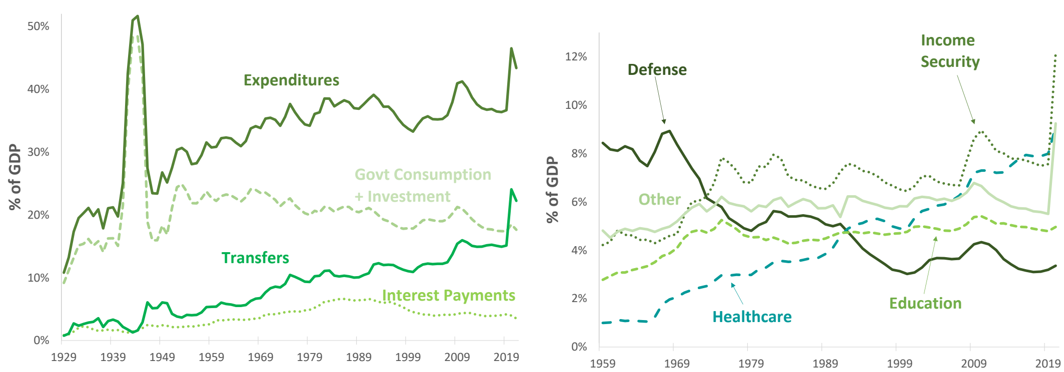

Figure 5: Expenditures by Category and Function, as percentages of GDP

The growth in the size of the U.S. government coincides with the creation (and expansion) of the Social Security system and the unemployment insurance program following the Great Depression, as well as the increase in public investment after WWII. This is illustrated by the left panel of Figure 5, which decomposes expenditures into the three sub-components

[1] See Table 1.1 here: https://www.cbo.gov/publication/58946.

[2] See Hall and Sargent [2020] and Blanchard [2023] for more on debt sustainability. When debt burdens become large, countries sometimes choose to default (as discussed in Chapter 22).

shown on the left-hand side of the government budget constraint in equation (4). Early on in the sample, government consumption and investment constituted the largest portion of expenditures and drove most of the trend. Over time, transfers expand significantly, overtaking G_t during the COVID-19 pandemic. In the data, transfers include both redistributive and insurance programs. Redistributive policies will be studied in Section 7.

The right panel of the figure shows the evolution of expenditures by function for selected items from 1959 and onward. Public education spending is relatively constant, accounting for 5 percent of GDP. Defense peaks in 1959 at 8 percent, but decreases significantly thereafter, to around 3 percent today. Health-care expenditure (including Medicare and Medicaid programs), on the other hand, show a steady rise, and now exceed 8 percent of GDP. Finally, "income security" spending, which includes unemployment insurance, retirement programs, disability and welfare, fluctuates around 7 percent of GDP. Spending on these items rises in recessions and decreases in booms, and as such these programs are typically referred to as "automatic stabilizers."

3 The effects of distortionary taxes

The objective of this section is to show how proportional taxes affect equilibrium allocations and prices. We do this in the context of the neoclassical growth model. Capital depreciates at rate δ. Households are infinitely-lived, discount at rate β, and enjoy utility each period from consumption and hours worked given by $u(c_t, \ell_t)$. They save in the form of capital, and rent capital and labor services to competitive firms at rates w_t and r_t. Firms produce according to a constant returns to scale production function $y_t = f(k_t, \ell_t)$. The government finances government consumption G_t and transfers T_t (which may be positive or negative) using proportional taxes on consumption, on labor income, and on rental income net of depreciation, τ_t^c, τ_t^ℓ, and τ_t^k. For now we assume no government debt (we introduce it in Section 4). The resource constraint is

$$C_t + G_t + K_{t+1} = f(K_t, L_t) + (1-\delta)K_t. \tag{3}$$

The government budget constraint is

$$G_t + T_t = \tau_t^c C_t + \tau_t^\ell w_t L_t + \tau_t^k (r_t - \delta)K_t. \tag{4}$$

The budget constraint for a representative household is

$$(1+\tau_t^c)c_t + k_{t+1} = (1-\tau_t^\ell)w_t \ell_t + k_t + (1-\tau_t^k)(r_t - \delta)k_t + T_t. \tag{5}$$

A government policy is a sequence $\{\tau_t^c, \tau_t^k, \tau_t^\ell, G_t, T_t\}_{t=0}^\infty$.

Definition 1 *: A competitive equilibrium given a policy* $\{\tau_t^c, \tau_t^k, \tau_t^\ell, G_t, T_t\}_{t=0}^\infty$ *is a sequence of allocations* $\{C_t, L_t, K_{t+1}\}_{t=0}^\infty$ *and prices* $\{w_t, r_t\}_{t=0}^\infty$ *such that* *i. Given policy and prices, the sequence* $\{C_t, L_t, K_{t+1}\}_{t=0}^{\infty}$ *maximizes household lifetime utility* $\sum_{t=0}^{\infty} \beta^t u(c_t, \ell_t)$ *subject to budget constraints of the form (5) for all* t*, initial capital* $k_0 = K_0$*, and a borrowing constraint* $k_{t+1} \geq 0\ \forall t$*.*

ii. The allocation $\{L_t, K_t\}_{t=0}^{\infty}$ *is a solution to the firm profit maximization problem at each date* t*, with* $\ell_t = L_t$ *and* $k_t = K_t$ *in equilibrium*

$$\max_{k_t, \ell_t} \{f(k_t, \ell_t) - w_t k_t - r_t k_t\}.$$

iii. The government budget constraint eq. (4) is satisfied at each date t*.*

At each date t, the first order conditions that define optimal saving and labor supply decisions for a household are

$$\begin{aligned} \frac{1+\tau_{t+1}^c}{1+\tau_t^c} \cdot \frac{u_c(c_t, \ell_t)}{u_c(c_{t+1}, \ell_{t+1})} &= \beta\left[1 + (1-\tau_{t+1}^k)(r_{t+1} - \delta)\right], \qquad (6) \\ -u_\ell(c_t, \ell_t) &= \frac{1-\tau_t^l}{1+\tau_t^c} \cdot w_t \cdot u_c(c_t, \ell_t). \end{aligned}$$

The conditions for profit maximization are

$$\begin{aligned} w_t &= f_\ell(k_t, \ell_t), \\ r_t &= f_k(k_t, \ell_t). \end{aligned}$$

In order to discuss the distortionary effects of taxation, it is useful to compute the Pareto optimal allocation. Because this is a representative agent economy, the efficient allocation can be found by solving the problem of a benevolent planner that maximizes lifetime utility subject to the resource constraint

$$C_t + G_t + K_{t+1} = f(K_t, L_t) + (1-\delta)K_t.$$

The first-order conditions to this problem are

$$\begin{aligned} \frac{u_c(C_t, L_t)}{u_c(C_{t+1}, L_{t+1})} &= \beta\left[1 + f_k(K_{t+1}, L_{t+1}) - \delta\right] \\ -u_\ell(C_t, L_t) &= f_\ell(K_t, L_t) \cdot u_c(C_t, L_t). \qquad (7) \end{aligned}$$

Comparing across these two sets of conditions, we can see how taxes change households incentives to save and to work. Absent taxes, households equate the inter-temporal marginal rate of substitution between consumption at t and at $t+1$ to one plus the marginal product of capital, net of depreciation. With taxes, households care instead about the gross after-tax return to saving, which is given by

$$\frac{1+\tau_t^c}{1+\tau_{t+1}^c}\left[1 + (1-\tau_{t+1}^k)(r_{t+1} - \delta)\right].$$

Holding r_{t+1} constant for a moment (in equilibrium r_{t+1} will depend on taxes) it is clear that taxes on rental income depress the after-tax return to saving, and that rising consumption taxes ($\tau_{t+1}^{c} > \tau_{t}^{c}$) work in the same direction. Similarly, taxes on labor income depress the return to working, as do taxes on consumption. Because taxes change workers' incentives to save and to work, they distort equilibrium allocations, and will typically reduce capital, labor supply, and output relative to the solution to the planner's problem.

3.1 Long-run distortions

Suppose we consider a steady state of the economy in which tax rates and allocations are constant. In such a steady state the household first-order conditions simplify to

$$\begin{aligned} 1 &= \beta \left[1 + (1-\tau^{k})(r-\delta)\right], & (8) \\ -u_{\ell}\left(c,\ell\right) &= \frac{1-\tau^{l}}{1+\tau^{c}} \cdot w \cdot u_{c}\left(c,\ell\right), & (9) \end{aligned}$$

where

$$\begin{aligned} w &= f_{\ell}(k,\ell), & (10) \\ r &= f_{k}(k,\ell). \end{aligned}$$

From the first of these it is immediate that a higher capital income tax τ^{k} must increase the steady state equilibrium rental rate for capital r. If the production function f has a Cobb-Douglas form, $f(k,\ell) = k^{\alpha}\ell^{1-\alpha}$, then $r = f_{k}(k,\ell) = \left(\frac{k}{\ell}\right)^{\alpha-1}$, and thus a higher rental rate corresponds to a lower capital-labor ratio. In particular,

$$\frac{k}{\ell} = \left(\frac{\alpha\left(1-\tau^{k}\right)}{\rho+\delta\left(1-\tau^{k}\right)}\right)^{\frac{1}{1-\alpha}} \tag{11}$$

where $\rho = \frac{1-\beta}{\beta}$ denotes the household's rate of time preference. Note that because a higher τ^{k} depresses the steady state capital-labor ratio, it will depress the steady state wage w, in addition to raising r. In contrast, labor and consumption taxes have no impact on the pre-tax prices r and w.

The effect of taxes on labor supply will depend on the preference specification, via the marginal utility terms in eq. (9). An increase in τ^{ℓ} affects labor supply in the current period directly by reducing the after-tax wage, inducing ℓ to fall via a substitution effect. At the same time, because the individual becomes poorer when labor income declines, consumption shrinks, which results in a higher marginal utility of consumption, incentivizing the agent to work more via an income effect. The total effect of an increase of labor income taxes on ℓ is therefore ambiguous, depending on the relative strength of substitution and income effects.

For an illustrative example we consider a particular utility function, made famous by Greenwood, Hercowitz, and Huffman (henceforth, GHH):

$$u(c,\ell) = \ln\left(c - \frac{\ell^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}}\right).$$

The GHH functional form is particularly tractable because consumption drops out of the the first-order condition for hours worked; thus, this utility function can be described as one in which there are no income effects.[3] Steady state labor supply is given by

$$\ell = \left(\frac{1-\tau^{\ell}}{1+\tau^{c}}\right)^{\phi} w^{\phi} \tag{12}$$

Note that hours worked depend only on the return to working, where ϕ defines the elasticity of hours to after-tax wages. Higher labor income or consumption taxes depress hours worked. Higher capital income taxes also depress hours, via their negative impact on w.

3.2 Tax incidence

In addition to simplifying the algebra, another feature of the GHH specification is that the steady state of the representative agent model specification is identical, at the aggregate level, to an alternative decentralization in which there are two household types: (i) workers, who rent labor services but own no capital, and (ii) capitalists, who own and rent out capital but who do not work. In what follows we focus on this worker-capitalist specification, because it allows for a discussion of tax incidence, namely the issue of who pays different sorts of taxes (in a representative agent setting, the representative household pays all taxes).[4]

From eqs. (10) and (12), we can solve for hours worked as a function of the capital to labor ratio

$$\ell = \left[\frac{1-\tau^{\ell}}{1+\tau^{c}}(1-\alpha)\left(\frac{k}{\ell}\right)^{\alpha}\right]^{\phi}$$

which, combined with eq. (11), gives an expression for steady state output

$$y = \left(\frac{k}{\ell}\right)^{\alpha}\ell = \left[\frac{1-\tau^{\ell}}{1+\tau^{c}}(1-\alpha)\right]^{\phi}\left[\frac{\alpha\left(1-\tau^{k}\right)}{\rho+\delta\left(1-\tau^{k}\right)}\right]^{\frac{\alpha(1+\phi)}{1-\alpha}} \tag{13}$$

Note that all three tax rates affect the level of steady state output, and that the level of output is decreasing in each tax rate. Capital income taxes depress the capital labor ratio,

[3] See Appendix 8.2 for an alternative utility function with income effects.

[4] Why is the steady state of the worker-capitalist model identical, given the GHH utility specification, to the steady state of the representative agent specification? The logic is that the level of consumption appears in neither the steady state first-order condition for saving, nor in the first order condition for hours worked. Thus the distribution of aggregate consumption between workers and capitalists has no impact on either steady state capital or steady state hours worked.

but their impact on output is amplified by the fact that a lower capital-labor ratio means lower wages, which in turn depress labor supply.

Consider the case with no transfers ($T = 0$). In steady state, workers consume

$$c^w = \frac{1-\tau^\ell}{1+\tau^c} w\ell = \frac{1-\tau^\ell}{1+\tau^c}(1-\alpha)y,$$

while capitalists consume

$$c^k = \frac{1-\tau^k}{1+\tau^c}(r-\delta)k = \frac{\rho}{1+\tau^c}\left[\frac{\alpha(1-\tau^k)}{\rho+(1-\tau^k)\delta}\right]y.$$

From these expressions, it is clear that labor taxes directly depress the consumption of workers, while capital income taxes directly depress the consumption of capitalists. Consumption taxes depress the consumption of both types. Note, however, that all three types of taxes indirectly depress the consumption of both types via their impact on equilibrium output.

If we are designing a tax system, we would like to know more about how effective different sorts of taxes are in terms of raising revenue, relative to how distortionary they are in terms of depressing output. To make further progress on this question in a tractable way, we now make two additional assumptions. First, we temporarily rule out consumption taxes by setting $\tau^c = 0$. Second, we assume that the government has to devote a fraction g of aggregate output to government purchases: $G = gY$. Thus, the steady state government budget constraint is

$$\begin{aligned} gy &= \tau^\ell wl + \tau^k(r-\delta)k \\ &= \tau^\ell(1-\alpha)y + \tau^k\rho\frac{\alpha}{\rho+(1-\tau^k)\delta}y \end{aligned}$$

From this budget constraint we can immediately solve for the locus of budget-balancing pairs (τ^ℓ, τ^k) :

$$\tau^\ell = \frac{1}{1-\alpha}\left[g - \tau^k\rho\frac{\alpha}{\rho+(1-\tau^k)\delta}\right] \tag{14}$$

Panel A of Figure 6 plots how the labor tax rate τ^ℓ varies with capital income tax rate τ^k according to eq. (14). The parameters used to construct the plot are $\alpha = 0.33$, $\delta = 0.07$, $\beta = 0.97$, $\phi = 1$ and $g = 0.2$. For each τ^k and the corresponding value for τ^ℓ given in Panel A, Panels B and C plot consumption of workers and capitalists, respectively. For comparison we also plot the levels of the two types' consumption for an economy without taxes. Note that as $\tau^k \to 1$, consumption of both types converges to zero. When τ^k is restricted to be non-negative, steady state consumption of capitalists is maximized at $\tau^k = 0$, while consumption of workers is a hump-shaped function of τ^k. Workers' utility depends on hours worked in addition to consumption, but given this utility function the consumption-equivalent argument of flow utility in equilibrium is proportional to consumption:

$$c^w - \frac{\ell^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} = \frac{1}{1+\phi}(1-\alpha)(1-\tau^\ell)y = \frac{1}{1+\phi}c^w.$$

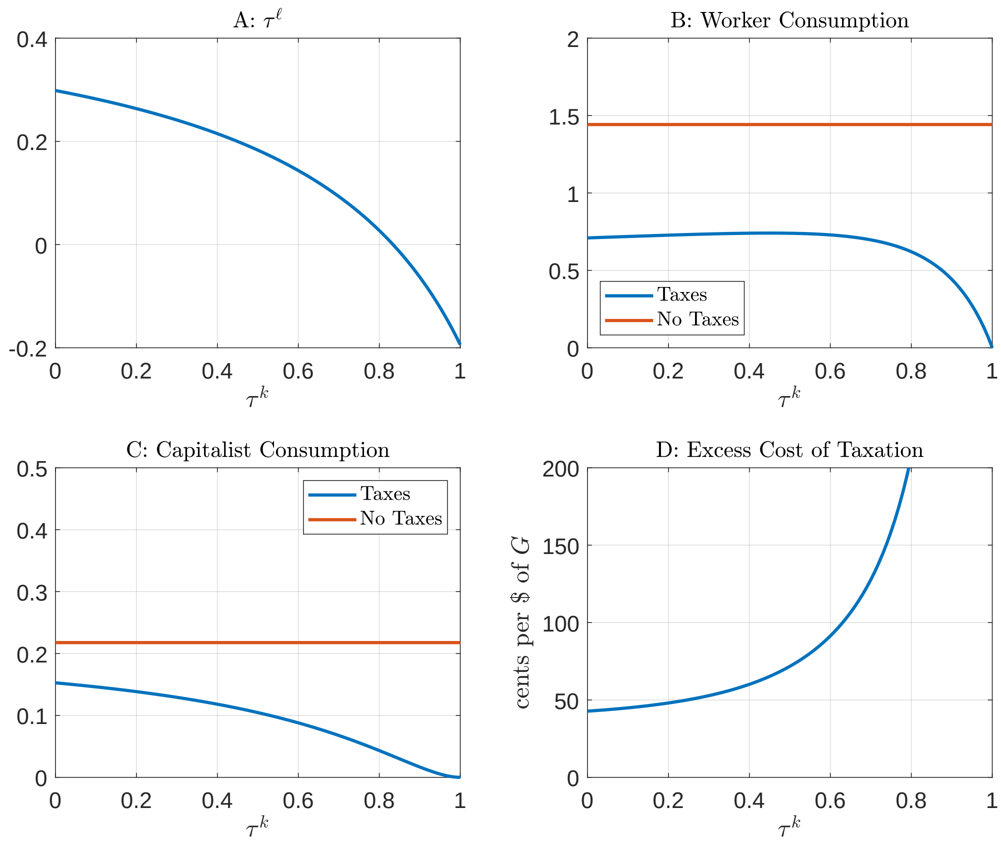

Figure 6: How Allocations Vary with τ^k

One way to define the deadweight cost of taxation is to ask by how much is total consumption reduced by taxes per unit of government consumption that the taxes finance. Let $c^w_{\tau=0}$ and $c^k_{\tau=0}$ denote the steady state consumption levels of workers and capitalists when $\tau^k = \tau^\ell = 0$, and define the excess cost of taxation as private consumption lost net of public consumption financed, measured per dollar of such spending:

$$\text{Excess Cost} = \frac{\frac{1}{1+\phi}\left(c^w_{\tau=0} - c^w\right) + c^k_{\tau=0} - c^k - gy}{gy}$$

If this ratio is equal to zero, then steady state utility in consumption units is reduced by one for each unit of government purchases. Panel D of Figure 6 plots the excess cost as τ^k varies (in the background τ^l is adjusted with τ^k to balance the government budget constraint). When $\tau^k = 0.2$, each dollar of government consumption comes with an excess cost of 50 cents, meaning that total private consumption is reduced by \$1.50 for each dollar of public goods provided. Clearly, the excess cost of taxation is always positive. Furthermore, the excess cost of taxation is increasing and convex in the capital tax rate τ^k.

3.3 Tax reform

Panel D of Figure 6 suggests that the excess cost of taxation is minimized when $\tau^k = 0$. The tax rate on capital income in the U.S., however, ranges from 10% to 37% (the ordinary income tax brackets in 2022, applied to capital held less than a year). From Panels B and C of the same figure, we see that if capital taxes were reduced to zero and labor taxes were raised to support the same $g = G/GDP$ ratio, then capitalists' consumption in steady state would be much higher without much reduction in workers' consumption. These findings might suggest that eliminating capital income taxes would be a good idea. However, the steady state associated with a lower τ^k has a larger capital stock, and if capital taxes are reduced it will take time for the economy to accumulate this extra capital. Additional capital accumulation will come at the cost of reduced current consumption. It is therefore important to analyze *transitional dynamics* in addition to steady states when considering tax reforms. We illustrate this with a simple example, using the worker-capitalist framework from the previous section and again assuming no transfers, $T_t = 0$.[5]

We start from a situation in which $\tau^k = 0.25$, a midpoint of the current tax brackets. Labor taxes are set to $\tau^l = 0.2529$, obtained from eq. (14) to sustain $g = 0.2$ given the parameters used in Section 3.2. We assume that the economy is in steady state until period 10, and that a switch to $\tau^k = 0$ is implemented, unexpectedly and permanently, in period 11. At that date, we increase the labor tax to $\tau^l = 0.2985$ so that eq. (14) still holds at $g = 0.2$ in the new steady state. We allow G_t to vary during transition to balance the government budget date by date given constant tax rates. While the initial and final steady states can be characterized analytically, the evolution of k_t during transition requires the use of computational methods. We assume that the economy has reached the final steady state by date T. Knowing the initial and final conditions, k_0 and k_T, respectively, all we need is a sequence $\{k_t\}_{t=1}^{T-1}$ consistent with the first-order conditions of workers and capitalists at dates $\{0, ..., T-2\}$. This is a system of $T-1$ inter-temporal first-order conditions and $T-1$ unknowns, which can be solved using a standard non-linear system of equations root-finding routine.[6]

[5]While the steady state of the representative agent model (RA) is identical to the worker-capitalist environment (WK), transitional dynamics are not the same. In the RA model, the first order condition with respect to capital includes the disutility of labor (due to non-separability between c and ℓ), whereas this is absent in the WK environment, since capitalists set $\ell = 0$. In the example computed above, the difference is numerically insignificant. However, it could be sizable with other preference specifications.

[6]The values for consumption in the inter-temporal first-order condition can be substituted out using the budget constraint of capitalists, $(1 + \tau_t^c)c^k + k_{t+1} = k_t + (1 - \tau_t^k)(r_t - \delta)k_t$. It is important to verify that T is large enough that the economy has indeed converged to the new steady state.

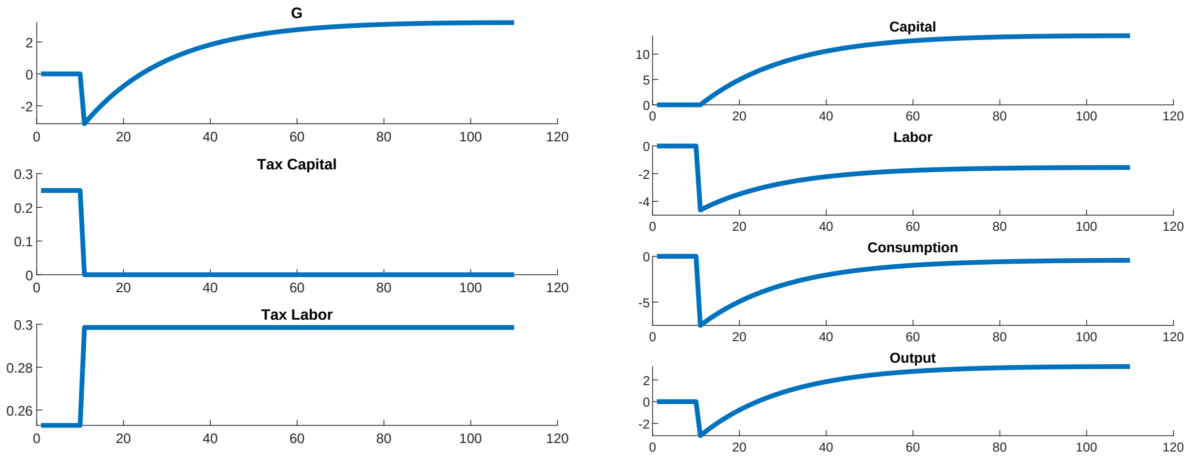

Figure 7: Eliminating capital income taxes. Except for the tax rates, variables are plotted as percentage deviations from their values in the pre-reform steady state.

The left panel of Figure 7 displays the evolution of capital and labor income taxes (exogenous parameters to the model), as well as the endogenous evolution of G_t. Public spending changes in response to the reform because prices and allocations change in equilibrium (as seen in the right panel), in turn affecting government revenues (recall that G/Y is identical in the initial and final steady states). The elimination of capital income taxes encourages capital accumulation, whereas the increase in labor income taxes discourages labor supply upon impact.[7] Over time, capital grows, and because this positively affects wages, labor supply gradually recovers, ending up only slightly below the initial steady state level. GDP tracks labor supply in the short run, and capital in the long run, declining right after the reform and recovering slowly over time. Aggregate consumption decreases initially, both because output is low, and because lower capital taxes are stimulating saving and investment. Subsequently, income rises and investment slows, pushing consumption back to a value close to the initial steady state.

While agents are eventually better off as the economy becomes more efficient (they work less than in the initial steady state but enjoy similar consumption), the exercise highlights that transitions can be painful. Whether the reform is beneficial for society overall depends on how much weight is assigned to capitalists versus workers. Capitalists are definitely better off, as their income always grows, whereas workers may be significantly worse off, as their tax burden increases (see, e.g., Domeij and Heathcote 2004.) Even when considering a representative household, whether public spending is valued or not is important in this calculation. The reform is more desirable if agents derive utility from government consumption – which

[7]The latter is an artifact of the specific utility function used, since it exhibits no income effects. In Appendix 8.2, we re-compute this experiment using an alternative specification.

rises over time – than if public spending is entirely wasteful.

3.4 The Laffer Curve

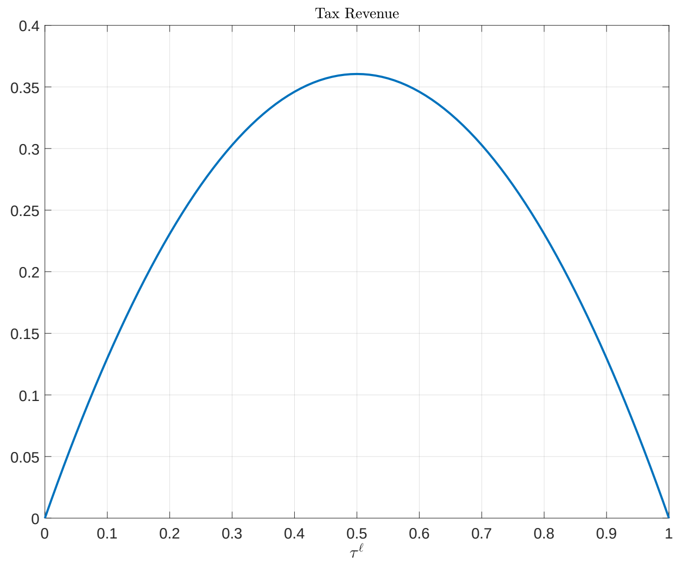

Figure 8: The Laffer curve

Another question of interest is: what combination of tax rates maximizes steady state tax revenue? A government that is fighting a war and wants to purchase the largest possible number of tanks might be especially interested in answering this question. When the planner only has access to taxes on labor earnings and rental income, and when tax rates must be positive, the revenue maximizing pair of tax rates is given by

$$\begin{aligned} \tau^{\ell} &= \frac{1}{1+\phi} \\ \tau^{k} &= 0. \end{aligned}$$

Note that the higher is the Frisch elasticity of labor supply, ϕ, the lower is the revenue-maximizing labor tax rate.[8]

Figure 8 plots tax revenue as a function of τ^{ℓ}, holding fixed τ^k at $\tau^k = 0$. This type of plot was popularized by Arthur Laffer in the 1970s, and is thus called a Laffer curve. Revenue is

[8] These rates are the solution to the problem:

$$\max_{\tau^{\ell} \geq 0, \tau^{k} \geq 0} \left\{ \tau^{\ell}(1-\alpha)y + \tau^{k}\rho \frac{\alpha}{\rho + (1-\tau^{k})\delta} y \right\},$$

hump-shaped in the tax rate, and for tax rates above the revenue-maximizing rate, raising rates reduces revenue, because the tax base shrinks faster than the rate increases. Such a situation is known as "being on the wrong side of the Laffer curve." To see why the Laffer curve must be hump-shaped, it is enough to observe that at $\tau^\ell = 0$, no revenue is raised because no taxes are levied, while at $\tau^\ell = 1$ no revenue is raised because hours worked and output are equal to zero, and thus there is nothing to tax.

3.5 Theories of G

So far, we have assumed that public spending is exogenously given and generates no benefits to society; revenues are "thrown into the ocean." However, as shown in Figure 5, public expenditures are composed of government consumption, public investment, transfers, and interest payments. Here, we briefly describe theories of government consumption and public investment. Interest payments will be described in the next section, after we introduce debt into the model, and transfers and welfare programs will be described at the end of the chapter.

First, let us focus on how to model government consumption. We typically assume that the government has the technology to provide public goods that are valued by society. These include defense, law and order, taking care of parks and common areas, sanitation, etc. A key assumption is that the government uses resources to produce these goods, which then provide utility to agents. We typically assume that agents derive utility from public and private goods, $u(c, g)$, with $u_g > 0$ and $u_{gg} \leq 0$. The first best solution (e.g., when the government has access to lump-sum taxation) prescribes equating the marginal utility of private consumption to the marginal utility of public consumption: $u_c(c, g) = u_g(c, g)$. When taxes are distortionary, the government must take into account, in addition, the deadweight losses associated with taxation.

A second strand of theories of government spending focuses on the role of the government to provide key infrastructure such as roads, bridges, and schools. Letting k_g denote 'public capital,' the neoclassical growth model augmented to include public investment involves a production function,

$$f(k, k_g, \ell) = Ak^\alpha \ell^{1-\alpha} k_g^\theta,$$

and a law of motion for k_g,

$$k_{g,t+1} = i_{g,t} + (1 - \delta_g)k_{g,t},$$

with $i_{g,t}$ denoting public investment in period t, and δ_g the rate of depreciation of public capital. The parameter θ controls the elasticity of output with respect to public capital. If

where output y is given by eq. (13). The first-order condition with respect to τ^ℓ gives the solution

$$\tau^\ell = \frac{1 - \tau^k \frac{\rho\phi}{1-\alpha} \cdot \frac{\alpha}{\rho + (1-\tau^k)\delta}}{1 + \phi}$$

Given this value for τ^ℓ, tax revenue is declining in τ^k at $\tau^k = 0$, indicating that $\tau^k = 0$ is the revenue-maximizing rate when tax rates must be positive.

$\theta > 0$, there are increasing returns to scale. Estimates of θ vary, from as low as 0.05 (see Leeper et al. [2010]) to as high as 0.39 (Aschauer [1989]). If revenue can be raised in a non-distortionary way, then it is optimal to equate the marginal products of private and public capital, net of depreciation, $f_{k,t+1} - \delta = f_{g,t+1} - \delta_g$.

4 Government debt and Ricardian Equivalence

We now add government debt to the analysis. To simplify the presentation, we abstract from capital, as this reduces the number of state variables. Using a two-period example, we present conditions under which debt is irrelevant. This result is known as "Ricardian Equivalence." We then introduce distortionary taxation and explain how the government can use debt to smooth out tax distortions over time. The optimal sequence of taxes solves what is known as the "Ramsey Problem."

Consider a two-period representative household model. The government can issue debt in period $t = 0$ and can levy lump-sum taxes or hand out lump-sum transfers in periods $t = 0$ and $t = 1$. Suppose that the government contemplates a lump-sum transfer T_0 in the first period, financed by issuing government debt B_0, which will be paid back by levying a lump-sum tax τ_1 in the second period.

The representative household has utility defined over consumption and hours worked in periods 0 and 1 given by

$$u(c_0, \ell_0) + \beta u(c_1, \ell_1).$$

Households choose labor supply, consumption and saving in each period, taking as given exogenous wages, w_0 and w_1, and an exogenous return to saving, r_0. Abstracting from other taxes, the budget constraints in the two periods are

$$\begin{aligned} c_0 + b_0 &= w_0 \ell_0 + T_0 \\ c_1 &= w_1 \ell_1 + (1 + r_1)\, b_0 - \tau_1, \end{aligned}$$

where b_0 is the amount of debt the household buys in the first period, and r_1 is the interest paid on that debt. Dividing the second equation through by $1 + r_1$ and adding it to the first equation expresses the budget constraint in present value form:

$$c_0 + \frac{c_1}{1 + r_1} = w_0 \ell_0 + \frac{w_1 \ell_1}{1 + r_1} + T_0 - \frac{\tau_1}{1 + r_1}.$$

Note that the value for debt purchases, b_0, drops out of the present-value version of the budget constraint. In addition, note that given a promised interest rate of r_1, the second period tax τ_1 will have to satisfy $\tau_1 = (1 + r_1)\, B_0 = (1 + r_1)\, T_0$. Thus, the tax and transfer terms in the present value budget constraint must sum to zero, regardless of the size of the initial transfer T_0. Since the tax and transfer scheme does not affect the lifetime budget constraint, the household's optimal allocation of consumption and hours must be identical to that in the case of zero taxes and transfers ($T_0 = \tau_1 = B_1 = 0$). The household must therefore respond to the initial transfer T_0 by increasing savings by exactly T_0. This extra savings will (i) exactly match the additional supply of government bonds issued, and (ii) provide exactly enough second period income to pay the expected lump-sum tax τ_1. This neutrality result is an example of *Ricardian Equivalence*. Note that this result hinges on the assumptions that taxes are lump-sum, that households face no credit constraints, and that the households who get the transfers are the same ones that must repay the debt. The result does not hinge on there being no capital. Using the same logic, it is possible to show that the result extends to an infinite-horizon economy (see also Barro 1974 and Heathcote 2005).

5 Ramsey Taxation

It is traditional in public finance to assume that the government cannot impose lump-sum taxes. Why? In a representative agent economy, there is no reason not to impose lump-sum taxes: such taxes are a distortion-free way to raise revenue. But in practice, actual households differ widely in terms of their income. Some households are so poor that they could not afford to pay a moderate lump-sum tax. Equally importantly, many people would find it unfair if the poor were expected to pay as much tax as the rich. Thus, the literature has focused on taxes that are *proportional* to income (see Chamley [1986] and Judd [1985] for early examples, and more recently Straub and Werning [2020]). Of course, one could consider making taxes a more complicated function of income – and we shall do so shortly – but proportional is simple, and simplicity can be viewed as a virtue.

But even if taxes are proportional to income, there is no need to tax different types of income at the same rate. In particular, earned income (income from labor) can be taxed at a different rate to unearned income (income generated by wealth). And if the government can save or borrow by issuing government debt, then it can also choose how tax rates should vary over time. We now consider a simple two period model and ask how a government that seeks to maximize the welfare of a representative agent should optimally set proportional taxes.

The timing assumptions are as follows. The government moves first, and announces tax rates for both periods, which we label periods 0 and period 1. To start, we assume that the government commits to these tax rates, and does not have the ability to deviate at $t = 1$ from the policies announced at $t = 0$. Later we will discuss how the analysis might change if the government does not have this commitment power.

A tax plan is *feasible* if there is a competitive equilibrium characterized by allocations and prices such that (i) those allocations are optimal choices for households and firms given prices and the tax rates described in the plan, (ii) the government budget constraint is satisfied, and (iii) markets clear.

A tax plan is *optimal* if it is feasible and the associated competitive equilibrium maximizes the welfare of the representative household. The optimal tax plan is called the *Ramsey plan*, and the associated equilibrium the *Ramsey equilibrium*.[9] In general, lots of different fiscal plans will be feasible, but only one will be optimal. There are two different approaches in the literature to solving for the Ramsey plan.

The first and more intuitive approach is to work with the full set of equilibrium equations and variables, and to conceptualize the planner choosing tax rates to maximize household welfare, internalizing how changes in tax rates will affect all equilibrium variables. This is called the dual approach.

An alternative approach, called the primal approach, is sometimes easier to implement (see Atkeson et al., 1999). Under the primal approach, we think of the planner as choosing equilibrium allocations for consumption and hours directly, subject to two sets of constraints. The first set of constraints ensure that allocations are technologically feasible. The second set of constraints ensure that there exists a set of tax rates such that the allocation is the competitive equilibrium given those taxes. These second constraints are called the *implementability constraints*. No equilibrium prices or tax rates appear in the primal problem. Once one has solved the primal problem, one can back out the tax rates that decentralize the solution in a final step.

5.1 The primal approach to optimal taxation: A simple example

The best way to understand how the primal approach works is to consider a simple example economy. Consider, in particular, the following two period model. There is a representative household with utility defined over consumption and hours worked in periods 0 and 1 given by

$$u(c_0, \ell_0) + \beta u(c_1, \ell_1).$$

The representative household supplies labor to a representative firm that produces output according to

$$\begin{aligned} y_0 &= A_0 \ell_0, \\ y_1 &= A_1 \ell_1, \end{aligned}$$

where A_t denotes potentially time-varying labor productivity. Because labor markets are assumed to be competitive, equilibrium wages equal productivities:

$$\begin{aligned} w_0 &= A_0, \\ w_1 &= A_1. \end{aligned} \tag{15}$$

Households (and the government) can save or borrow using a storage technology that converts one unit of output at date 0 into $1 + r$ units of output at date 1, where r is an

[9] After Frank Ramsey, who wrote a handful of important papers in economics and more in the fields of mathematics and philosophy, before his death in 1930 at the age of 26.

exogenous constant.[10] Let b_{t-1} denote household wealth at the start of period t. Assume that $b_{-1} = 0$.

The government must finance exogenous expenditures g_0 and g_1 in periods 0 and 1. It can raise tax revenue via proportional taxes on labor income at rates τ_0 and τ_1, and by taxing income from wealth in period 1 at rate τ_1^b.[11] A *tax plan* is a vector $\left\{\tau_0, \tau_1, \tau_1^b\right\}$.

Let b_0^g denote government savings in the storage technology at date 0. The government budget constraints for periods 0 and 1 are

$$\begin{aligned} g_0 + b_0^g &= \tau_0 w_0 \ell_0, \\ g_1 &= \tau_1 w_1 \ell_1 + \tau_1^b r b_0 + (1+r) b_0^g. \end{aligned}$$

Note that government savings at $t = 0$ delivers income at $t = 1$. If $b_0^g < 0$, the government is borrowing. These two constraints can be combined to give

$$g_0 + \frac{g_1}{1+r} = \tau_0 w_0 \ell_0 + \frac{\tau_1 w_1 \ell_1}{1+r} + \frac{\tau_1^b r b_0}{1+r}.$$

Given taxes and wages, the representative household solves

$$\begin{aligned} & \max_{\{c_0, \ell_0, c_1, \ell_1, b_0\}} u(c_0, \ell_0) + \beta u(c_1, \ell_1) \\ & s.t. \\ c_0 &= (1-\tau_0) w_0 \ell_0 - b_0, \\ c_1 &= (1-\tau_1) w_1 \ell_1 + (1 + r(1-\tau_1^b)) b_0, \end{aligned}$$

where again the two budget constraints can be collapsed to give

$$c_0 + \frac{c_1}{1 + r(1-\tau_1^b)} = (1-\tau_0) w_0 \ell_0 + \frac{(1-\tau_1) w_1 \ell_1}{1 + r(1-\tau_1^b)}.$$

The first-order conditions that characterize the solution to the household's problem are

$$\begin{aligned} u_{c,0} w_0 (1-\tau_0) &= -u_{\ell,0}, \\ u_{c,1} w_1 (1-\tau_1) &= -u_{\ell,1}, \\ u_{c,0} &= \beta(1 + r(1-\tau_1^b)) u_{c,1}, \end{aligned} \tag{16}$$

where $u_{c,t}$ denotes the marginal utility of consumption in period t.

[10] One interpretation might be that the economy is small and open, and r is the world interest rate.

[11] The household has no wealth at date 0, so a tax on wealth or income from wealth at date 0 would not raise any revenue. If the household did have wealth at date 0, the government would like to tax that wealth, since that would effectively amount to a non-distortionary lump-sum tax. In the spirit of not allowing for lump-sum taxation, it is typically assumed that the Ramsey planner cannot tax initial wealth, or that there is an upper bound on the feasible initial tax rate.

Resource feasibility in this economy can be summarized by a single equation, which states that the present value of private plus public consumption is equal to the present value of output

$$c_0 + g_0 + \frac{c_1 + g_1}{1+r} = A_0\ell_0 + \frac{A_1\ell_1}{1+r}. \tag{17}$$

What about the implementability constraints? An allocation is a competitive equilibrium if it satisfies the three first-order conditions from the household problem, the two equilibrium expressions for wages, and the household and government budget constraints. We now show that these six equations can be collapsed into a single implementability condition. The idea is to take the household lifetime budget constraint, and to use the household first-order conditions to substitute out for $(1-\tau_0)w_0$, $(1-\tau_1)w_1$, and $1+r(1-\tau_1^b)$. In particular, the first-order condition for saving implies that, in any competitive equilibrium, it must be the case that

$$1 + r(1-\tau_1^b) = \frac{u_{c,0}}{\beta u_{c,1}},$$

while those for labor supply imply

$$(1-\tau_t)w_t = -\frac{u_{\ell,t}}{u_{c,t}}.$$

After these substitutions, the household lifetime budget constraint can be written as

$$c_0 + \frac{c_1}{\frac{u_{c,0}}{\beta u_{c,1}}} = -\frac{u_{\ell,0}}{u_{c,0}}\ell_0 - \frac{u_{\ell,1}}{u_{c,1}}\ell_1\frac{1}{\frac{u_{c,0}}{\beta u_{c,1}}}, \tag{18}$$

or, after multiplying through by $u_{c,0}$, as

$$u_{c,0}c_0 + \beta u_{c,1}c_1 = -u_{\ell,0}\ell_0 - \beta u_{\ell,1}\ell_1.$$

This is the implementability condition. It can be shown that any allocation that satisfies the resource constraint and the implementability condition can be implemented by some feasible tax plan (see, e.g., Atkeson et al., 1999).

It should be clear that the implementability condition embeds the household optimality conditions for labor supply and savings in addition to the household budget constraint. But what about the government budget constraint? We do not have to worry separately about that, because if the resource constraint and the household budget constraint are satisfied, the government budget constraint must also be satisfied, by Walras Law.

The *Ramsey problem* is to maximize lifetime utility for the representative agent, subject to the resource and implementability constraints. Writing this problem as a Lagrangian, with multipliers λ and μ on the resource and implementability constraints, we have

$$\begin{aligned}
&\max_{c_0,c_1,\ell_0,\ell_1} \{u(c_0,\ell_0) + \beta u(c_1,\ell_1) \\
&+\lambda\left(A_0\ell_0 + \frac{A_1\ell_1}{1+r} - c_0 - g_0 - \frac{c_1+g_1}{1+r}\right) \\
&+\mu\left(u_{c,0}c_0 + \beta u_{c,1}c_1 + u_{\ell,0}\ell_0 + \beta u_{\ell,1}\ell_1\right)\}
\end{aligned}$$

The first-order conditions are

$$\begin{aligned}
u_{c,0} - \lambda + \mu u_{c,0} + \mu u_{cc,0} c_0 + \mu u_{\ell c,0} \ell_0 &= 0, \\
u_{\ell,0} + \lambda A_0 + \mu u_{\ell,0} + \mu u_{c\ell,0} c_0 + \mu u_{\ell\ell,0} \ell_0 &= 0, \\
\beta u_{c,1} - \frac{\lambda}{1+r} + \beta\mu u_{c,1} + \beta\mu u_{cc,1} c_1 + \beta\mu u_{\ell c,1} \ell_1 &= 0, \\
\beta u_{\ell,1} + \frac{1}{1+r}\lambda A_1 + \beta\mu u_{\ell,1} + \beta\mu u_{c\ell,1} c_1 + \beta\mu u_{\ell\ell,1} \ell_1 &= 0.
\end{aligned}$$

where, for example, $u_{cc,0}$ denotes $\partial u_{c,0}/\partial c_0$.

These four first-order conditions alongside eqs. (17) and (18) constitute six equations that can be used to solve for the six unknowns $(c_0, c_1, \ell_0, \ell_1, \lambda, \mu)$. Thus, one can solve for the Ramsey allocation. Consider, in particular, a separable utility function of the form

$$u(c, \ell) = \frac{c^{1-\gamma}}{1-\gamma} - \frac{\ell^{1+\sigma}}{1+\sigma}.$$

In this case, the cross derivative terms drop out, and the second derivatives simplify to

$$u_{cc,t} c_t = -\gamma u_{c,t}, \; u_{\ell\ell,t} \ell_t = \sigma u_{\ell,t}.$$

Thus, the first-order conditions can be written as

$$\begin{aligned}
u_{c,0}(1 + \mu - \mu\gamma) &= \lambda \\
u_{\ell,0}(1 + \mu + \mu\sigma) &= -\lambda A_0 \\
\beta u_{c,1}(1 + \mu - \mu\gamma) &= \frac{1}{1+r}\lambda \\
\beta u_{\ell,1}(1 + \mu + \mu\sigma) &= -\frac{1}{1+r}\lambda A_1
\end{aligned}$$

Comparing the first and the third, it is immediate that, at an optimum

$$\beta(1+r) u_{c,1} = u_{c,0}.$$

Comparing the first and the second, we see that

$$u_{c,0} A_0 \cdot \frac{1 + \mu - \mu\gamma}{1 + \mu + \mu\sigma} = -u_{\ell,0}$$

Similarly, the third and the fourth give

$$u_{c,1} A_1 \cdot \frac{1 + \mu - \mu\gamma}{1 + \mu + \mu\sigma} = -u_{\ell,1}$$

Comparing these expressions to the first order conditions for saving and for working in the original economy (eqs. 16), and noting that, in equilibrium, $w_0 = A_0$ and $w_1 = A_1$, it is clear that the only way both sets of first order conditions can be satisfied at the same allocation is if

$$\begin{aligned} \tau_1^b &= 0, \\ \tau_0 &= \tau_1 = 1 - \frac{1+\mu-\mu\gamma}{1+\mu+\mu\sigma} = \frac{\mu(\sigma+\gamma)}{1+\mu+\mu\sigma}. \end{aligned}$$

Thus, this simple example illustrates two classic results in the Ramsey taxation literature. First, the government should commit to neither tax nor subsidize income from savings (see Chamley 1986 and Judd 1985). Second, the labor tax rate should be constant over time, and will be positive as long as either g_0 or g_1 is strictly positive (so that revenue must be raised). This result is described as *tax smoothing*, and the idea is that because distortions from taxes increase with the tax rate in a convex fashion, constant tax rates are preferable to time-varying tax rates (see Barro [1979] and Lucas and Stokey [1983]).

5.2 Time consistency

Let us assume that parameters are such that the household saves in period 0 under the Ramsey plan. One parameter configuration that would deliver this is $\beta(1+r) = 1$ – so that the Ramsey allocation features $c_1 = c_1$ – and $A_0 > A_1$, so the Ramsey allocation features $A_0\ell_0 > A_1\ell_1$.

Recall that our analysis above presumed that the government announced a tax plan at date 0 and stuck to the plan at date 1. Suppose now that we give the planner the ability to redesign taxes in period 1. At $t = 1$ the planner can raise revenue either by taxing labor earnings (as promised in the original plan) or by taxing household income from savings. What combination of taxes would a benevolent planner choose? The answer is that such a planner would set $\tau_1 = 0$ and τ_1^b as high as necessary to fund required government purchases. The reason is that once period 1 rolls around, household wealth b_1 is already determined, and taxes on income from wealth are effectively a lump-sum tax. In contrast, taxes on labor earnings are distortionary. Because the planner would like to deviate from the Ramsey plan at date 1, given the chance to do so, the plan is said to be *time inconsistent* (see Kydland and Prescott [1977]).

There are many policy questions where time consistency arises as a central issue. For example, Chapter 22 focuses on debt policy. If governments can commit to repaying debt, they will be able to borrow cheaply. But those promises to repay may not be time consistent, in the sense that once debt has been accrued the government might be better off defaulting.

Does the fact that the Ramsey plan described above is time inconsistent mean that we should not take it too seriously as a practical policy prescription? It is certainly useful to know what policy would be optimal given commitment and to understand how the planner might be tempted to deviate from the Ramsey plan. It might be possible to design institutions in such a way that the planner has more commitment power – for example, by writing a constitution that precludes frequent tax changes. At the same time, there is a large literature that attempts to characterize the best time consistent policies (see Klein et al. [2008]).

6 Debt and pensions with overlapping generations

We close this chapter by studying fiscal policy in a non-dynastic economy and return to the two-period overlapping-generations endowment economy discussed in Chapter 5.5. We extend this model to incorporate government debt, a pay-as-you-go (PAYG) pension system, and taxes. To simplify the exposition, we assume a small open economy with access to a global bond market.

The population grows at a constant rate n. Let $N_t = (1+n)^t$ denote the size of the newborn young population at date t. The share of young people is $N_t / (N_t + N_{t-1}) = (1+n) / (2+n)$. Only the young work and their endowment of efficiency units grows at rate γ. Thus, the labor income of an individual born in period t is $y_t = (1+\gamma)^t \omega$ and aggregate labor income is $Y_t = y_t N_t = (1+n)^t (1+\gamma)^t \omega$.

The government operates a PAYG pension system and provides a public good G_t. The pension system pays p_t to every old individual, financed by taxing labor income at rate τ_p. The pension per retiree is therefore

$$p_t = (1+n)\, \tau_p y_t.$$

Spending on the public good is assumed to be a fixed fraction of GDP: $G_t = gY_t$. Government spending is financed by taxing labor income by a flat tax τ_t and by issuing debt. We abstract from taxes on capital income. The government budget constraint is given by

$$G_t + p_t N_{t-1} + (1+r)\, B_{t-1} = (\tau_p + \tau_t)\, Y_t + B_t,$$

where B_t is the issuance of new debt that matures next period and the interest rate r is exogenous and fixed. Because the pension system is self-financing, the budget can be expressed as follows,

$$g + \frac{1+r}{(1+n)\,(1+\gamma)} b_{t-1} = \tau_t + b_t,$$

where b_t is the debt to GDP ratio at the end of period t.

Individuals maximize discounted utility, $u(c_{y,t}) + \beta u(c_{o,t+1})$, subject to budget constraints when young and old,

$$\begin{aligned} c_{y,t} + a_t &= (1 - \tau_t - \tau_p)\, y_t, &(19)\\ c_{o,t+1} &= (1+r)\, a_t + p_{t+1}, &(20) \end{aligned}$$

where a_t denotes saving at t, and where, in equilibrium, $p_{t+1} = (1+n)(1+\gamma)\tau_p y_t$. The sequence for optimal consumption can then be computed from two equations: a lifetime budget constraint and an Euler equation:

$$c_{y,t} + \frac{c_{o,t+1}}{1+r} = \left[1 - \tau_t - \tau_p + \frac{(1+n)(1+\gamma)}{1+r}\tau_p\right] y_t \tag{21}$$

$$1 = (1+r)\beta \frac{u'(c_{o,t+1})}{u'(c_{y,t})}. \tag{22}$$

Note, first, that a pension system is equivalent to issuing a particular form of government debt. To see this, consider an individual who has no pension tax or transfer ($\tau_p = p_{t+1} = 0$) but who is forced to purchase $b_{p,t}$ government bonds with a promised return r_p. The budget constraints for this individual would be

$$c_{y,t} + a_t + b_{p,t} = (1 - \tau_t) y_t, \tag{23}$$

$$c_{o,t+1} = (1+r) a_t + (1+r_p) b_{p,t}. \tag{24}$$

If we set $b_{p,t} = \tau_p y_t$ and $1+r_p = (1+n)(1+\gamma) \approx (1+n+\gamma)$, then the budget constraints with "pension debt" (23-24) are equivalent to those with a pension system, (19-20). Note that the return on forced saving in the PAYG pension system is equal to the growth rate of output. Therefore, the pension system increases the present value of household income for the young if and only if wage growth exceeds the interest rate, i.e., iff $\gamma + n > r$ (this condition determines whether the right-hand side of eq. (21) is increasing in τ_p). This insight has an important implication: if the interest rate is larger than the growth rate of output $n + \gamma$ (a "normal" scenario) then the pension system is effectively a tax on the young generation. The generation who are old when the system is first introduced gain because they receive benefits without having paid taxes when they were young. But the current young generation and all future generations lose because they get a higher return on private savings than on pension contributions. However, if the interest rate is *lower* than the wage growth rate $n+\gamma$, then all generations gain from introducing a pension – both the initial old generation and all generations of young. In this case, introducing a PAYG pension system is Pareto improving. This corresponds to an equilibrium featuring dynamic inefficiency, as discussed in Chapter 6.4.1.

A major change in this overlapping-generations model relative to the standard infinite-horizon model is that Ricardian equivalence no longer holds. In particular, the timing of taxes and transfers now matters for the distribution of consumption across cohorts, and for the trajectory of aggregate consumption. To see this, consider a one-time transitory tax holiday in period t in an economy with zero initial debt. Thus, $\tau_t = 0$ and $B_t = gY_t$. Moreover, assume that the government finances the repayment of this debt by increasing taxes in period $t+1$ and does not issue debt thereafter. Note, first, that this "tax holiday" increases the present value of consumption for generation t and lowers the present value of consumption for generation $t+1$; see equation (21). Thus, the debt-financed tax cut shifts the tax burden from generation t to generation $t+1$. The result is that aggregate consumption will increase in period t and fall in period $t+2$.[12] Thus, Ricardian equivalence breaks down. This is different from the case we studied in Section 4. There, Ricardian equivalence held because the government's debt policy did not affect the present value of taxes for the representative household. Here, in contrast, debt policy reshuffles the tax burden across generations.

This analysis has a number of implications for actual policies. Real-world pension systems are not purely PAYG and many countries have accumulated pension funds. However, public pension savings are relatively small: the U.S. Social Security Administration is scheduled to deplete its trust fund by 2033.

Two factors have strained public pension systems in OECD countries. First, the number of retirees relative to the number of workers has increased and will continue to increase in coming decades. Population aging is driven by both lower mortality (retirees living longer) and by lower fertility. This can be interpreted as a lower n in the model above. Second, the productivity growth rate has fallen in recent decades (secular stagnation). For example, the U.S. growth rate of GDP per capita – γ in our model – fell from 2.3% between 1950 and 2000 to 1.2% between 2000 and 2020.

The analysis above suggests that pension promises are a form of debt. What is the total level of effective government debt for the U.S. federal government? A narrow definition of debt corresponds to the value of government bonds outstanding. For the U.S., federal debt held by the public was 94% of GDP in the second quarter of 2023.[13] A more comprehensive definition includes the implicit debt in the pension system, i.e., the present value of future federal pension promises. For the U.S. federal government this measure of debt is \$65.9 trillion, or almost 2.5 times annual GDP.[14] This massive figure excludes the future costs of Medicare (the federal health care program for retirees). These two measures of government debt—94% versus 94+245=339%—are strikingly different. The *de facto* debt burden therefore depends on how seriously one should take promises about financial debt versus promises of future pension benefits. An outright default on nominal debt is ruled out by the U.S. Constitution. However, the government could increase surprise inflation and thereby inflate away some of the debt. Pension promises, in contrast, do not enjoy any constitutional protection and the government is always free to reduce social security benefits or raise the age at which people are eligible to collect them.

7 Taxes and transfers as instruments for redistribution

An important function of government is to redistribute and provide social insurance. To this end, the tax and transfer system includes a wide array of taxes, social insurance programs

[12] The effect on aggregate consumption in $t+1$, $C_{t+1} = c_{0,t+1}N_t + c_{y,t+1}N_{t+1}$, is ambiguous because the tax cut will increase $c_{0,t+1}$ and lower $c_{y,t+1}$.

[13] Debt held by the public excludes the holdings of government debt by Federal government entities such as the Social Security Trust Fund, but includes debt held by the Federal Reserve.

[14] Source: 2023 OASDI Trustees Report, Table VI.F2.

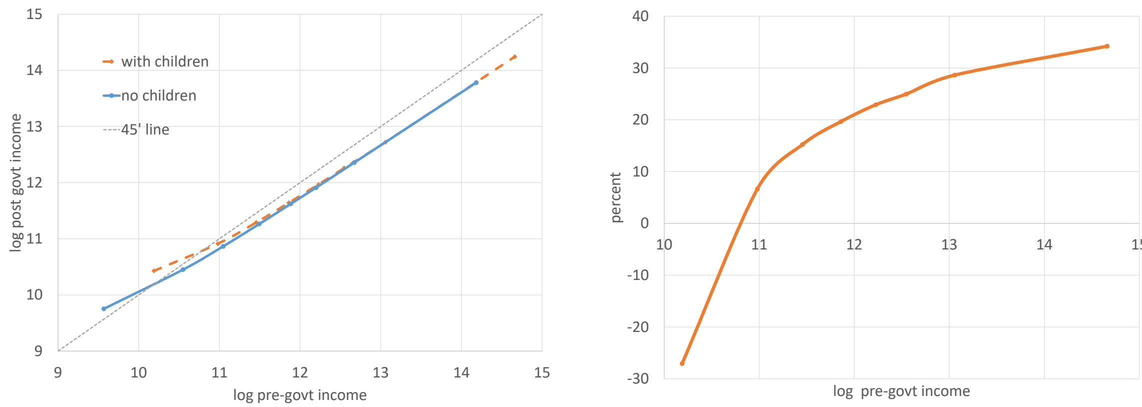

Figure 9: Left plot: Scatter plot of pre-government income against and post-government income for percentiles of U.S. households. Right plot: Average net tax rates by household income, defined as taxes minus transfers as a share of income, for households with children. Source: Congressional Budget Office (CBO), 2016.

and means-tested benefits at different levels of government (federal, state, and local). To illustrate the extent of redistribution embedded in the U.S. system, Figure 9 plots pre- versus post-government income for each percentile of the pre-government income distribution. Pre-government income is income before taxes and transfers. Post-government income is disposable income, defined as pre-government income plus transfers minus taxes. Each dot in the plot shows average pre- and post-government income for one percentile of the pre-government household income distribution. The relationship is approximately linear, except at the lowest income percentiles. This suggests that the U.S. tax- and transfer system can be well approximated by a log-linear function:

$$y - T(y) \quad = \quad \lambda y^{1-\tau}, \tag{25}$$

where y is pre-government household income, $T(y)$ is taxes minus transfers, and $y - T(y)$ is disposable income. The parameter λ controls the level of taxation, while the parameter τ can be interpreted as a measure of tax progressivity. To see this, note that when $0 < \tau < 1$, the tax system is *progressive* in the sense that the marginal tax rate $T'(y)$ is larger than the average tax rate $T(y)/y$ for any positive income level. Conversely, when $\tau < 0$, the marginal tax rate is lower than the average tax rate, $T'(y) < T(y)/y$, implying that taxes are *regressive*. When $\tau = 0$, the tax system is flat, with a constant marginal tax rate $T'(y) = T(y)/y = 1-\lambda$.

The right panel of Figure 9 plots average net tax rates, defined as taxes minus transfers divided by pre-government income. The picture illustrates that the average net tax rate is increasing with income. The U.S. tax and transfer system can therefore be said to be *progressive*.

7.1 A macro model of progressivity

We now illustrate the effects of tax progressivity on inequality and the macro economy using a simple static model of redistribution. The economy is populated by a unit continuum of individuals indexed by i. The utility function u is

$$u\left(c_i, \ell_i, G\right) = \log c_i - \frac{\ell_i^{1+\sigma}}{1+\sigma},$$

where c_i and ℓ_i are consumption and labor supply of individual i.

The tax and transfer system is assumed to take the log-linear form described in equation (25). The government's budget must be balanced, which imposes a constraint on the set of feasible fiscal policy choices (τ, λ, G).

The aggregate resource constraint dictates that output is spent on either private consumption or on public goods:

$$Y = \int_0^1 c_i \, di + G.$$

Individuals differ with respect to labor productivity. Their labor income is $w_i \ell_i$, where w_i is individual i's productivity. Individuals have no wealth, so consumption must equal disposable income:

$$c_i = \lambda \left(w_i \ell_i\right)^{1-\tau}.$$

Taking a first-order condition with respect to hours worked, one can solve in closed form for the equilibrium allocation. Hours worked and consumption are given by

$$\log \ell_i = \frac{\log(1-\tau)}{1+\sigma}, \tag{26}$$

$$\log c_i = \log \lambda + (1-\tau)\frac{\log(1-\tau)}{1+\sigma} + (1-\tau)\log w_i. \tag{27}$$

Hours worked are falling in τ but are independent of the individual wage, w_i. Progressivity ($\tau > 0$) reduces hours because workers internalize that if they increase hours they will face a higher marginal tax rate, depressing the after-tax return. In the limit as $\tau \to 1$, workers anticipate that disposable income will equal λ, irrespective of hours worked, and thus hours will shrink to zero. Hours are independent of the wage because the utility function is in the balanced growth class.

Consumption is increasing in individual productivity, w_i. Tax progressivity dampens the pass-through from wages to consumption: a one percent increase in wages translates to a $1-\tau$ percent increase in consumption. Thus, tax progressivity reduces consumption inequality: the variance of log pre-government earnings is $var(\log w)$, while the variance of log consumption is $(1-\tau)^2 \, var(\log w)$. In conclusion, this simple economy illustrates the fundamental trade-off between efficiency and redistribution in a setting with an empirically plausible tax- and transfer system: higher progressivity reduces hours worked (lower efficiency) but also reduces consumption inequality (more redistribution). For more discussion of this trade-off, see Heathcote et al. [2017].

References

David A. Aschauer. Is public expenditure productive? *Journal of Monetary Economics*, 23 (2):177–200, 1989.

Andrew Atkeson, V.V. Chari, and Patrick Kehoe. Taxing capital income: A bad idea. Technical Report 2331, Federal Reserve Bank of Minneapolis, 1999. URL https://www.minneapolisfed.org/research/quarterly-review/taxing-capital-income-a-bad-idea.

Robert J. Barro. Are government bonds net wealth? *Journal of Political Economy*, 82 (6):1095–1117, 1974. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/1830663.

Robert J. Barro. On the determination of public debt. *Journal of Political Economy*, 82(6): 1095–1117, 1979.

Olivier Blanchard. *Fiscal Policy under Low Interest Rates*. The MIT Press, 2023.

Christophe Chamley. Optimal taxation of capital income in general equilibrium with infinite lives. *Econometrica*, 54(3):607–622, 1986. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1911310.

David Domeij and Jonathan Heathcote. On the distributional effects of reducing capital taxes. *International Economic Review*, 45(2):523–554, 2004.

George J. Hall and Thomas J. Sargent. Debt and taxes in eight U.S. wars and two insurrections. Working Paper 27115, National Bureau of Economic Research, May 2020. URL http://www.nber.org/papers/w27115.

Jonathan Heathcote. Fiscal policy with heterogeneous agents and incomplete markets. *Review of Economic Studies*, 72:161–188, 2005.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics*, 132(4):1693–1754, 2017.

Kenneth L. Judd. Redistributive taxation in a simple perfect foresight model. *Journal of Public Economics*, 28(1):59–83, 1985. ISSN 0047-2727. doi: https://doi.org/10.1016/0047-2727(85)90020-9. URL https://www.sciencedirect.com/science/article/pii/0047272785900209.

Paul Klein, Per Krusell, and José-Víctor Ríos-Rull. Time-Consistent Public Policy. *The Review of Economic Studies*, 75(3):789–808, 07 2008. ISSN 0034-6527. doi: 10.1111/j.1467-937X.2008.00491.x. URL https://doi.org/10.1111/j.1467-937X.2008.00491.x.

Finn E. Kydland and Edward C. Prescott. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491, 1977. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/1830193.

Eric M. Leeper, Todd B. Walker, and Shu-Chun Yang. Government investment and fiscal stimulus. *Journal of Monetary Economics*, 57:253–92, 2010.

Robert E. Lucas and Nancy L. Stokey. Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics*, 12(1):55–93, 1983. ISSN 0304-3932. doi: https://doi.org/10.1016/0304-3932(83)90049-1. URL https://www.sciencedirect.com/science/article/pii/0304393283900491.

Ludwig Straub and Iván Werning. Positive long-run capital taxation: Chamley-judd revisited. *American Economic Review*, 110(1):86–119, January 2020. doi: 10.1257/aer.20150210. URL https://www.aeaweb.org/articles?id=10.1257/aer.20150210.

8 Appendix to Government and public policy (Marina Azzimonti, Jonathan Heathcote, Kjetil Storesletten)

8.1 Data Appendix

The data on revenues, expenditures, and deficits is obtained from the Bureau of Economic Statistics (BEA), Table 3.1 "Government Current Receipts and Expenditures," which is part of the NIPA tables[15]. The series, expressed in current billion dollars, span the interval 1929-2021. They include Federal, State, and Local government budget measures (sometimes referred to as General Government or National Government statistics). In the plots, the series are expressed as percentages of "GDP," corresponding to *Gross Domestic Product* (line 1 of Table 1.5.5).

Figure 1: Data is obtained from the OECD Dataset: National Accounts at a Glance, and the variable corresponds to *Total expenditure of general government, percentage of GDP*.

Figure 2: "Revenues" correspond to *Total Receipts* (line 34 of Table 3.1) and "Expenditures" to *Total Expenditures* (line 37 of Table 3.1). We use total rather than current measures because these include public investment.

[15] See https://apps.bea.gov/iTable

Figure 3: All series are obtained from Table 3.1 "Government Current Receipts and Expenditures," constructed by the BEA. "Income Taxes" corresponds to *Personal current taxes* (line 3), "Sales and Import Taxes" to *Taxes on production and imports* (line 4), "Corporate Taxes" to *Taxes on corporate income* (line 12), and "Social Ins. Taxes" to *Contributions for government social insurance* (line 7).

Figure 4 - Left Panel: "Total Deficit or Surplus," in Figure 4, is constructed as the difference between Revenues and Expenditures (defined above),

$$\text{Total Deficit=Revenues-Expenditures.}$$

"Net Interest" is the difference between *Interest and Miscellaneous Receipts* (line 11 of Table 3.1) and *Interest Payments* (line 27 of Table 3.1). The government simultaneously owns assets that yield interest and owes debt for which it has to pay interest. In the figure, we plot net interest payments. The "Primary Deficit" is defined as

$$\text{Primary Deficit= Total Deficit - Net Interest.}$$

Figure 4 - Right Panel: FRED provides debt series for Federal and State governments between 1946 and 2021. "Federal Debt" corresponds to *Federal Government; Debt Securities and Loans; Liability, Level* (or FGSDODNS), while "State and Local Debt" corresponds to *State and Local Governments; Debt Securities and Loans; Liability, Level* (or SLGSDODNS). Both series are obtained from the Flow of Funds tables constructed by the Board of Governors of the Federal Reserve Bank System. It is worth noticing that the Federal Debt series does not correspond exactly to the one provided by the White House historical series due to differences in accounting methods (i.e. which items are included and timing in which certain transactions are incorporated when computing the flow of funds). The series between 1916 and 1945 are obtained from the Survey of Current Business, September 1946 page 13, Table 5. They correspond to Net Public Debt, end of calendar year.

Figure 5 - Left Panel: All series are obtained from Table 3.1 (described above). "Govt Consumption (+ Investment)" is the sum of *Consumption expenditures* (line 20) and *Gross government investment* (line 39). "Transfers" is the sum of *Current transfer payments* (line 22) and *Subsidies* (line 30). "Interest Payments" are gross, obtained from line 27 (i.e. we are not including interest receipts). The sum of these is equal to Expenditures, defined above.

$$\text{Expenditures= Govt Consumption (+ Investment) + Transfers + Interest Payments.}$$

Figure 5 - Right Panel: All series are obtained from Table 3.16 "Government Current Expenditures by Function" constructed by the BEA. "Defense" corresponds to *National Defense* (line 7), "Healthcare" corresponds to *Health* (line 28), and "Education" is obtained

directly form line 30. "Income Security" is obtained from line 36. It includes *Disability* (line 37), *Welfare and social services* (line 39), *Unemployment* (line 40), *Retirement* (line 38) and other income insurance programs (line 41). "Other" is constructed as the sum of *General public service* (line 2), *Public order and safety* (line 8), *Economic affairs* (line 13), *Housing and community services* (line 27), *Recreation and culture* (line 29), minus *Interest payments* (line 5).

8.2 Tax reform with wealth effects

In this section, we re-compute the tax reform from Section 3.3, but assuming that utility takes the form

$$u(c,\ell) = \ln c - \frac{\ell^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}}.$$

The first order condition with respect to labor implies

$$\ell^{1/\phi} = \left(\frac{1-\tau_t^l}{1+\tau_t^c}\right) w \frac{1}{c_w} \quad \text{with} \quad c_w = \left(\frac{1-\tau_t^l}{1+\tau_t^c}\right) w.$$

The labor supply is *independent* of taxes in this case, $\ell = 1$. This happens because the substitution effect, that would make ℓ decline when after-tax labor income goes down is exactly offset by the income effect, caused by a decline in c_w that results in lower after-tax income.

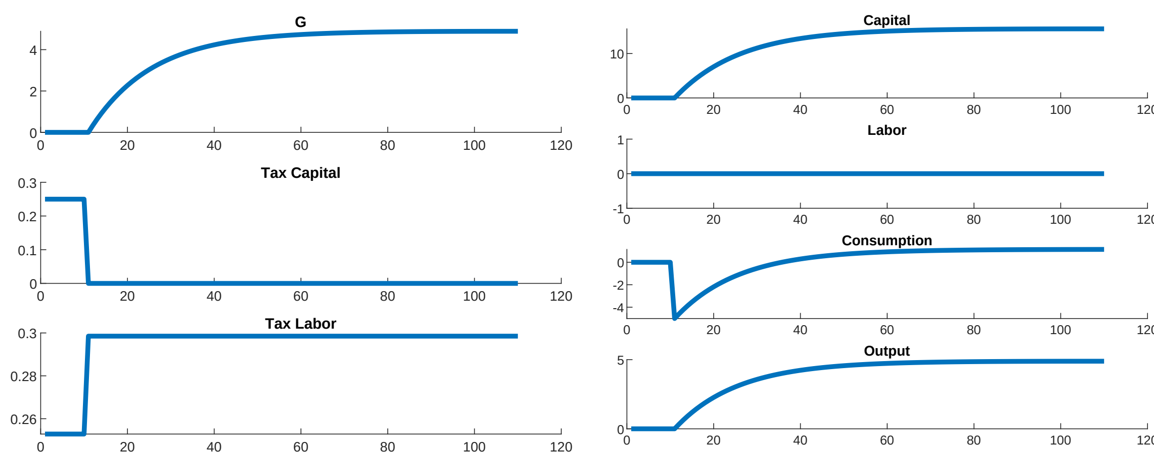

Figure 10: Eliminating capital income taxes (with wealth effects)

The main difference between Figure 7 and Figure 10 is that now labor supply remains constant when labor taxes are increased. As a result, we do not observe a decline in output, which allows the government to have higher G (recall that the exercise is constructed such that $G/Y = 0.2$ throughout the simulation). In the long-run, because labor does not go down, there is higher GDP and aggregate consumption is slightly higher. The tax reform is more effective in this scenario because the costs of replacing capital taxes with labor taxes are smaller when wealth effects are present.