Pre-announced optimal tax reform

David Domeij[a,b], Paul Klein[*,a]

[a] *Institute for International Economic Studies, Stockholm University, 106 91 Stockholm, Sweden*

[b] *Department of Economics, Northwestern University, 2003 Sheridan Road, Evanston, IL 60208-2600, USA*

First draft: 12 January 1998

This version: 25 February 1998

Abstract

Optimal tax policies in dynamic models have unappealing features. In particular, optimal tax reform typically involves a large initial accumulation of government assets which is responsible for a large part of the welfare gains from optimal tax reform. In this paper, we investigate the robustness of these findings by studying optimal tax policy in a standard growth model when a reform has to be announced in advance of its implementation. We find that this requirement leads to an optimal solution which is considerably more reasonable than optimal tax reforms studied previously. Using numerical calculations, we find that the optimal pre-announced tax reform involves only a small initial accumulation of government assets. We also find that the welfare gains from optimal tax reform are reduced by no more than a third when the government is required to pre-announce reform about 14 years in advance, and that this reduction is mainly due to the delay itself rather than the effect of pre-announcement on the character of the optimal tax reform. This leaves us with a welfare gain corresponding to an increase in consumption of about 1 percent from a tax reform with reasonable properties.

Keywords: Optimal taxation, commitment, frontloading.
JEL classification: E62, H21.

1 Introduction

Optimal tax policies in dynamic models typically exhibit some rather disturbing properties. The optimal tax scheme in Turnovsky & Brock (1980) has the government collecting all its revenues in the first period, setting all subsequent tax rates to zero. In Chari, Christiano & Kehoe (1994) (henceforth CCK), the optimal tax

[*] Correspondence to: klein@iies.su.se or domeijd@iies.su.se.

We thank Torsten Persson and Kjetil Storesletten for useful comments and suggestions and Harald Uhlig for inspiring us to write this paper. Generous financial support from *Tom Hedelius och Jan Wallanders stiftelse för samhällsvetenskaplig forskning* and *Carl Mannerfelts fond* is gratefully acknowledged.

reform involves an enormous initial-period consumption boom in spite of the preferences for consumption smoothing of the model's households. Jones, Manuelli & Rossi (1997) (henceforth JMR) show that, in a large class of models, optimal tax rates are all equal to zero in a steady state, calling this a "disturbing but essential feature of the optimal tax code".

In this paper, we present an economic environment where optimal tax reform looks quite reasonable. The setup involves only a small modification of the standard one, namely that the government has to announce reform in advance of its implementation. We argue that this modification makes sense in terms of the spirit of the Ramsey (1927) optimal taxation exercise. Our approach, we claim, takes to its logical conclusion the idea that a model of optimal taxation should be set up in such a way as to make the government take the full incentive effects of taxation into account when setting tax rates.

The existing literature on Ramsey optimal taxation in dynamic models (see, for example, Chamley (1981), Judd (1985) and CCK) takes as given the idea that lump-sum taxation is not feasible. Although the informational restrictions that lie behind this fact are not usually spelt out in detail, an underlying idea is that the government should be modelled as taking the incentive effects of taxation into account when it sets tax policy. For this to be consistent with rationality on the part of the government, tax policy must be set before households make their choices. Only then is it rational for the government to internalize the effects of its policy on the expectations of the households. Consequently, the existing literature on Ramsey taxation assumes that the government is able to precommit in the initial period to the (possibly state contingent) sequence of taxes from the initial period onwards.

In this paper, we argue that this amount of precommitment is still not enough; that it still leaves room for something that looks a lot like lump-sum taxation. Of course it is widely recognized that allowing the government freely to set the capital income tax in the initial period is quite literally to allow lump-sum taxation. Indeed, Turnovsky & Brock (1980) call it "introducing a lump sum tax through the back door". Consequently it is usually assumed that the initial period capital income tax rate is exogenously fixed. But we claim that even this is not enough. The intuitive reason is as follows. When the initial period capital income tax rate is fixed, it is true that the government cannot decide on the size of the fraction of initial holdings taken from the public. But it can influence its *value* by manipulating the time path of labor income taxes and expected future capital income taxes. In particular, what it can do is to stimulate a consumption boom in the initial period, which brings down the shadow value of the initial-period capital stock. Indeed, this is precisely what numerical exercises with Ramsey taxation in dynamic models have shown; see for example CCK. Moreover, the solution is typically characterized by a very high (expected) capital income tax rate in the period following the initial period. We claim that this property of the solution is contrary to the spirit of the Ramsey equilibrium. The reason is that the government engineers the boom only because it ignores the disincentive effects the low initial shadow value of capital would have had on pre-reform accumulation of capital *if it had been known beforehand*.

In an attempt to avoid an extreme initial consumption boom, JMR impose various restrictions on tax rates and bond issues. In this paper, we try to achieve the same result by instead attacking what we claim is the essential source of the initial consumption boom property of the Ramsey optimal tax reform, i.e. the insufficient amount of precommitment on the part of the government. We increase the degree of precommitment by separating the date of announcement from the date of implementation of reform. More specifically, we force the government to announce the optimal policy reform T periods in advance and investigate what happens to the optimal tax code as T changes, and, in particular, as T becomes a very large but finite number.[1] In this way, we force the government to take the disincentive effects of taxation on *pre-reform behavior* into account.

The main results are as follows. When $T \geq 2$, the initial consumption boom disappears and is replaced by a dip. As T is increased further, the size of this dip diminishes so that consumption eventually comes close to increasing monotonically over time towards the new steady state. As $T \to \infty$, capital accumulation is carried out mostly before the implementation of reform, and the limiting path comes close to exhibiting an immediate jump at the time of implementation ($t = 0$) to the new steady state. To support this equilibrium, the capital income tax in period 0 shrinks from 214 percent at the traditional solution ($T = 0$) to about -34 percent at $T \gg 0$. Meanwhile, the degree of frontloading[2] diminishes. At $T = 0$, about 9 percent of public spending is financed by returns on government assets in the steady state. This figure falls to about 2 percent as $T \to \infty$.

[1] We will sometimes express this by saying that $T \to \infty$. Of course, this does not imply that $T = \infty$ which means that there is no reform at all.

[2] For the purposes of this paper, a tax policy is said to be frontloaded if part of government spending on the balanced growth path is financed by returns on assets accumulated during the transition. For details, see section 6.

Concerning the sequence of solution paths created by changing T, we find that this sequence of paths eventually converges to a limiting path as $T \to \infty$, but rather slowly. While the solution does become much more reasonable when there is a pre-announcement period of about four years, T has to increase to about 20 years before further increases in T have only a small quantitative impact.

In terms of welfare, we find that the gains from optimal tax reform are reduced by as little as a third when the government is required to pre-announce reform about 14 years in advance, leaving a welfare gain corresponding to an increase in consumption of about 1 percent from a reform with quite reasonable properties. We also find that this reduction in the welfare gain is mainly due to the delay itself (postponing a good thing is a bad thing) rather than the effect of pre-announcement on the character of the optimal tax reform. When we control for the pure delay effect, we find that the welfare gains are more or less unchanged by the requirement of pre-announcement.

The remainder of the paper is organized as follows. Section 2 describes the economic environment. Section 3 discusses optimal revenue-neutral tax reform with and without pre-announcement. Section 4 describes how the model's parameters are calibrated to fit the post-war experience of the United States. Section 5 describes the numerical methods we have used. Section 6 presents the results for the benchmark case. In section 7, we investigate the robustness of the results in the benchmark case by studying some extensions. Section 8 concludes.

2 The economic environment

The economic environment is a deterministic neoclassical growth model with variable leisure and a government sector. We will denote consumption by c_t, hours worked as a fraction of total available time by h_t, the capital stock by k_t, the real pre-tax interest rate by r_t, the pre-tax wage rate by w_t, the stock of government bonds by b_t, and the price of government bonds by q_t. In the tradition of Ramsey optimal taxation, we impose linear tax schedules, denoting the flat tax rates on labor and capital income by τ_t^h and τ_t^k respectively. Government spending consists of purchases g_t and transfers T_t. As final piece of notation, partial derivatives will be written in the following way. Let $g : \mathbb{R}^n \to \mathbb{R}$ and let $i \in \{1, 2, ..., n\}$. Then

$$g_i(x) := \frac{\partial g(x)}{\partial x_i}. \tag{1}$$

When x is indexed by t, we will sometimes write

$$g_{i,t} := g_i(x_t). \tag{2}$$

The economy has a representative household with the period utility function $U(c,h)$ and a competitive single production sector with the aggregate production function $f(k,h)$.

Each household is assumed to be of measure zero, and the representative household therefore takes market prices as parametrically given. It solves

$$\max_{\langle c_t, k_t\rangle_{t=0}^{\infty}} \left[\sum_{t=0}^{\infty} \beta^t U(c_t, h_t)\right] \tag{3}$$

$$\text{s.t.} \begin{cases} c_t + k_{t+1} + q_t b_{t+1} = \left(1-\tau_t^h\right) w_t h_t + \left[1 + \left(1-\tau_t^k\right)(r_t - \delta)\right] k_t + b_t + T_t;\ t = 0, 1, ... \\ k_0 \text{ given} \\ b_0 \text{ given} \\ \liminf_{t\to\infty} \left\{\beta^t \lambda_t k_{t+1}\right\} \geq 0 \\ \liminf_{t\to\infty} \left\{\left(\prod_{k=0}^{t} q_k\right) b_{t+1}\right\} \geq 0 \end{cases} \tag{4}$$

where $\lambda_t = U_1(c_t, h_t)$ is treated as parametrically given by the household.[3] Meanwhile, a representative price-taking firm maximizes profits so that factors of production are paid their marginal products according to

$$\begin{cases} r_t = f_1(k_t, h_t) \\ w_t = f_2(k_t, h_t). \end{cases} \tag{5}$$

For convenience, we define the after-tax return on capital via

$$R_t := 1 + \left(1-\tau_t^k\right)(r_t - \delta).$$

The consumer's and firm's problem are solved in the context of an economy which faces the aggregate resource constraint

$$c_t + k_{t+1} + g_t = (1-\delta) k_t + f(k_t, h_t). \tag{6}$$

This description of our model economy ignores the objectives and constraints of the government, which is the topic of the next section. Instead we end this section by

[3] Another way of putting this is to say that $\lambda_t = U_1(C_t, H_t)$ where X_t is the population average of x_t.

defining a competitive equilibrium *given* sequences of government policies $\langle b_t \rangle_{t=0}^{\infty}$, $\langle g_t \rangle_{t=0}^{\infty}$, $\langle \tau_t^h \rangle_{t=0}^{\infty}$ and $\langle \tau_t^h \rangle_{t=0}^{\infty}$. It is an allocation $\langle c_t \rangle_{t=0}^{\infty}$, $\langle h_t \rangle_{t=0}^{\infty}$, $\langle k_t \rangle_{t=0}^{\infty}$ and prices $\langle r_t \rangle_{t=0}^{\infty}$, $\langle w_t \rangle_{t=0}^{\infty}$, $\langle q_t \rangle_{t=0}^{\infty}$ such that the allocation satisfies the aggregate resource constraint and maximizes household utility and firm profits given the prices, which must also satisfy (5). Bond prices q_t are determined by the implicit requirement that bond markets clear in every period. We will assume that the household's problem is such that, once government policies are given, the following equations, along with (5), uniquely characterize the competitive equilibrium.

$$\begin{cases} U_{1,t} = \beta R_{t+1} U_{1,t+1};\ t = 0, 1, ... \\ q_t U_{1,t} = \beta U_{1,t+1};\ t = 0, 1, ... \\ U_{2,t} + U_{1,t} \left(1 - \tau_t^h\right) w_t = 0;\ t = 0, 1, ... \\ c_t + k_{t+1} + q_t b_{t+1} = \left(1 - \tau_t^h\right) w_t h_t + R_t k_t + b_t + T_t;\ t = 0, 1, ... \\ c_t + k_{t+1} + g_t = f\left(k_t, h_t\right);\ t = 0, 1, ... \\ k_0,\ b_0 \text{ given} \\ \lim_{t\to\infty} \left\{\beta^t U_{1,t}\left(c_t, h_t\right) k_{t+1}\right\} = 0 \\ \lim_{t\to\infty} \left\{ \left(\prod_{k=0}^{t} q_k\right) b_{t+1} \right\} = 0. \end{cases} \tag{7}$$

3 Pre-announced optimal tax reform

3.1 Standard Ramsey optimal reform

The standard Ramsey equilibrium is that competitive equilibrium allocation which delivers the maximum value of the household's objective function subject to the constraints

$$\begin{cases} \langle g_t \rangle_{t=0}^{\infty} \text{ given} \\ \langle T_t \rangle_{t=0}^{\infty} \text{ given} \\ g_t + T_t + b_t = \tau_t^h w_t h_t + \tau_t^k \left(r_t - \delta\right) k_t + q_t b_{t+1};\ t = 0, 1, ... \\ b_0 \text{ given} \end{cases} \tag{8}$$

A useful theorem in this context is that the constraints (7) and (8) faced by the government can be summarized in a single equation, namely the so-called *implementability* condition. It says that

$$\sum_{t=0}^{\infty} \beta^t \left[U_{1,t}\left[c_t - T_t\right] + U_{2,t} h_t\right] = U_{1,0}\left[b_0 + R_0 k_0\right]. \tag{9}$$

To prove this theorem, one can follow the approach of Lucas & Stokey (1983). Intuitively speaking, the implementability condition summarizes three properties of any feasible competitive equilibrium allocation, namely (i) the government's budget flow budget constraint is satisfied in each period, (ii) the household's Euler equations are satisfied in each period, and, finally, (iii) the present value of outstanding government debt converges to zero. Formally, then, the Ramsey problem is to maximize

$$\mathcal{U} = \sum_{t=0}^{\infty} U\left(c_t, h_t\right)$$

subject to (6) and (9).

Unless we impose further constraints on the problem, however, the Ramsey problem has a rather trivial solution; it consists in setting τ_0^k high enough to generate revenue that suffices to buy enough government claims on the private sector so that taxes need never be levied again. This is in fact always possible even if k_0 is arbitrarily small, since the household can borrow from the government the amount it needs to cover any taxes due for payment in period 0. There is common agreement that this solution goes against the spirit of the Ramsey problem, which after all is premised upon the idea that non-distortionary (lump-sum) taxation is not feasible. Since the capital levy τ_0^k is unanticipated by the household, it works just like a lump sum tax, it is typically constrained in most analyses of Ramsey optimal taxation in dynamic models (see e.g. CCK).

In fact, it is common to go one step further and constrain not just the initial capital levy but the entire sequence of tax rates (see e.g. JMR) by requiring it to be within a certain band, say between 0 and 100 percent. The motivation for this restriction is typically that the solution would otherwise be "unreasonable" in some (not always clearly articulated) sense. One sense in which it might be thought to be unreasonable is simply that the equilibrium allocation looks funny. It involves a sudden jump in consumption to a level way above the steady state generated by average current tax policy, and then an equally sudden dive to a point below the old steady state. Consumption then slowly recovers and settles down somewhere strictly above the old steady state level. In this paper we argue that it is suspect in a deeper sense; in fact in the same sense as the initial-capital-levy solution described in the previous paragraph. It is suspect because it goes against the spirit of the Ramsey optimal taxation in that it allows for something that is conceptually akin to lump-sum taxation.

This claim is based on the reason for why an initial consumption boom is a property of the Ramsey equilibrium with τ_0^k fixed. The point is that anything which can bring down the value of $U_{1,0}[b_0 + R_0 k_0]$ is a good thing from the point of view of the government. This term represents the market value of initial assets held by the private sector. An initial capital levy τ_0^k brings this down in a direct way. If τ_0^k is fixed exogenously, it is optimal to bring the market value of initial assets down in an indirect way, viz. by bring down the value of the market price of capital, which in equilibrium is $U_{1,0}$. The period utility function U being concave in consumption, $U_{1,0}$ is brought down by making c_0 a large number. Of course, the government cannot control c_0 directly, but it can manipulate the rate of return between period 0 and period 1 in such a way as to encourage consumption in the initial period. This is done by setting $\tau_1^k \gg 0$ and $\tau_0^h \leq 0$.

We argue, then, that both the initial capital levy phenomenon and the initial consumption boom property are inconsistent with the spirit of the Ramsey optimal taxation exercise. We also argue that they should be remedied in the same way. Fixing tax rates exogenously or forcing them into arbitrary bands, moreover, does not seem like an appropriate way of proceeding. If we don't like the solution to an optimization problem, the only honest option is to modify the objective function and/or the constraints in a way that we can motivate independently of whether the solution as such is funny or not. In the next section, we present such a modification.

3.2 Pre-announced optimal reform

3.2.1 The intuitive idea

In this paper, we modify the constraints of the Ramsey optimal taxation exercise in a way that we claim is in line with the original spirit of the Ramsey optimal taxation idea (as formulated originally by Pigou). Ramsey's (1927) setting was of course static, so that all agents were aware of all the taxes when they made their decisions about tax bases. No taxes were unanticipated. Generalizing to the stochastic case, this would not necessarily mean that taxes could not deviate from their conditional means. What it would mean is that the conditional mean of each tax rate is known and taken into account by households when the relevant decisions concerning the formation of the tax base are made.

Generalizing to a dynamic setting is a more sensitive matter. To stay true to the Ramsey spirit, the conditional means of all the relevant tax rates should be known at the time of formation of the respective tax bases. Presumably this should apply to the capital stock as well as for any other tax base. But the decisions that determined the capital stock at the time of reform are all the savings decisions from the dawn of time ($t = -\infty$) until that time. Hence to enable agents to choose the initial capital stock taking all the relevant tax rates into account, tax policy has to be *pre-announced*. In principle, it should be pre-announced arbitrarily far in advance.

3.2.2 The mathematical formalization

From the perspective of the optimizing government, pre-announcement takes the form of tax rates being constrained to their old values (given by the current sub-optimal policy) in periods $t = 0, 1, ..., T-1$. The reform is then implemented in period T. Denote the old (exogenously given) value of τ_t^h by $\bar{\tau}^h$ and the old value of τ_t^k by $\bar{\tau}^k$. Using the household's optimum conditions, the pre-announcement constraints become

$$\begin{cases} U_{2,t} + U_{1,t}\left(1-\bar{\tau}^h\right) w_t = 0; & t = 0, 1, ..., T-1 \\ U_{1,t} - \beta\left[1 + \left(1-\bar{\tau}^k\right)\left(r_{t+1} - \delta\right)\right] U_{1,t+1} = 0; & t = 0, 1, ..., T-2. \end{cases}$$

The second of these constraints is of a non-standard kind, as discussed by Marcet & Marimon (1995). As a consequence, the Lagrange multiplier η_t associated with the capital income tax constraint becomes a state variable in the sense that its initial value is given ($\eta_{-1} = 0$) and that the value of η_{t-1} is relevant for decisions taken in period t.

For any fixed T, $\tau_t^k = \bar{\tau}^k$ for $t = 0, 1, 2, ..., T-1$. To avoid an initial confiscation, $\tau_0^k = \bar{\tau}^k$ even when $T = 0$. Thus when $T > 0$, the first freely chosen capital income tax is τ_T^k, and when $T = 0$, the first freely chosen capital income tax rate is $\tau_1^k = \tau_{T+1}^k$. Regardless of the value of T, the first freely chosen labor income tax is τ_T^h.

4 Model calibration

4.1 Current government policy

4.1.1 Taxes

We set the labor tax rate $\tau^h = 0.26$ and the capital income tax rate $\tau^k = 0.44$. These numbers are in the same region as those reported by McGrattan, Rogerson

& Wright (1994) and Mendoza, Razin & Tesar (1994).

4.1.2 Spending

Government purchases g_t are set constant at a level so that the steady state ratio of government purchases to GDP generated by the model with current policy is 0.19, which is the average post-war figure for the United States. Transfer payments are also constant and such as to balance the budget in the steady state generated by the model with current policy. This implies a steady state ratio of transfers to GDP of 0.046. This is slightly less than the U.S. post-war average, which is about 10 percent. The reason for this discrepancy is that we omit some sources of tax revenue, such as sales taxes, from the model.

4.1.3 Debt

Initial government debt b_0 is set to zero in the benchmark case. See section 7 for a discussion of what happens when this assumption is modified.

4.2 Preferences and technology

The production function is defined via

$$f(k,h) = k^{\theta} h^{1-\theta} \tag{10}$$

where $\theta = 0.36$ to match the post-war capital share of income in the United States. The period utility function is defined via

$$U(c,h) = \alpha \ln c + (1-\alpha) \ln (1-h)$$

where $\alpha = 0.2757$ so as to deliver $h = 0.23$ in the steady state generated by the model with current government policy. This value is consistent with the post-war data from the United States, provided that one interprets total available time as 24 hours per day.

The value of k_0 is set so as to equal the steady state capital stock generated model with current government policy.

Each period is taken to represent two years, which we think of as the minimal time it can realistically take between a tax reform being proposed and its being implemented. Consequently we set $\delta = 0.2$ and $\beta = 0.92$. See section 7 for a discussion of what happens when these parameters are calibrated so that each period represents a single year.

5 Numerical methods

Our first attack on the problem was to use a version of the parameterized expectations algorithm (PEA) described in Marcet & Marshall (1994), modified to take into account the time-dependence of the problem.[4] This approach turned out to be extremely slow and was therefore abandoned. Instead we proceed as follows. We begin by dividing the problem into a pre-reform time-dependent part ($t < T$) and a post-reform stationary part ($t \geq T$). Our algorithm starts by guessing k_T and λ, the shadow price associated with the implementability condition. We then solve for the post-reform part of the problem, i.e. we use the sufficient conditions for maximizing

$$\sum_{t=T}^{\infty} \beta^{t-T} U\left(c_t, h_t\right) \tag{11}$$

subject to the relevant constraints, in particular that k_T is given. The solution for this part is found using a version of the parameterized expectations algorithm (PEA) described in Marcet & Marshall (1994). Retaining the guess of k_T and letting k_{T+1} and h_T be given from the solution to the post-reform problem, we proceed to solve the pre-reform problem, i.e. we maximize

$$\sum_{t=0}^{T} \beta^{t} U\left(c_t, h_t\right)$$

subject to the relevant constraints. Since h_T is given, we ignore the first order condition with respect to it. However, we retain the first order condition for the optimal choice of k_T in spite of the fact that it is fixed. To avoid the resulting system of equations from being overdetermined, we let k_0 yield. We then have an exactly determined finite-dimensional non-linear system of equations. To solve this system, we first reduce its dimensionality by using the resource constraint to solve for consumption. Then we exploit the linearity in the Lagrange multipliers to solve for them by solving a system of linear equations. This leaves us with a $2T-$dimensional system of equations in $k_0, k_1, ..., k_{T-1}$ and $h_0, h_1, ..., h_{T-1}$, which we solve using a numerical minimization algorithm.

At this stage, we check whether k_T and λ were guessed correctly in the sense that the solution for k_0 is equal to its exogenously given value and that the implementability condition is satisfied. If not, we update k_T and λ until convergence.

[4] The problem is time-dependent between $t = 0$ and $t = T - 1$, since the number of periods remaining until $T-$day $(T - t)$ is a relevant state variable.

6 Results

In this section, we present the positive and welfare properties of the solution for different values of T. We discuss the policies and allocations separately, but also the relationship between them.

6.1 Positive properties of the solution

6.1.1 Policies

Closely related to the initial consumption boom and an initial high anticipated capital income tax rate, we have the phenomenon of *frontloading*. For the purposes of this paper, a tax policy is said to be frontloaded if part of government spending in the steady state is financed by returns on assets accumulated during the transition. This phenomenon has been discussed recently by JMR who call it a "disturbing but essential feature of the optimal tax code." Indeed, as JMR show, there is a wide class of models in which the optimal tax code is characterized by *complete* frontloading in the sense that all taxes tend to zero.[5] While JMR may be right in saying that frontloading is an essential feature of Ramsey optimal reform, we find that the extent of frontloading depends negatively on the size of the interval between announcement and implementation of the reform. Indeed, in our setup, it is quantitatively insignificant when policy is pre-announced sufficiently far in advance.

As exhibited in Table 1 and Figure 3, the quantitative importance of frontloading diminishes as T increases, although it does not disappear entirely. When $T = 0$, as much as 9 percent of government expenditures are covered by returns on government assets in the steady state. When $T = 15$, this figure is only 2 percent. Similarly, as T increases, the initial accumulation of assets diminishes and the steady state government assets converge to a very small positive number.

Corresponding to the smaller initial accumulation of assets, the first freely chosen capital income tax rate falls as T increases. Indeed, it is negative for $T > 6$, and converges to about -34 percent as T becomes large. This provides the incentives for the ever-faster convergence of the capital stock to its new steady state level as T increases. For details of the optimal capital income taxes, see Table 3 and Figure 4.

[5] This class of models have labor services supplied jointly with human capital and constant returns in the accumulation of human and physical capital. Our benchmark model does not feature human capital accumulation and hence is not a member of this class.

Similarly, the increase in τ_T^h from 0.8 percent when $T = 0$ to more than 30 percent as soon as $T > 0$ corresponds to the disappearance of the initial period consumption boom when $T > 0$. Meanwhile, the dynamics of labor tax rates are trivial; as seen in Table 3, they converge almost immediately to their new steady state value when $T > 0$.

6.1.2 Allocations

When $T = 0$ we have the standard case as studied by CCK. Figure 2 exhibits the initial period consumption boom and Figure 3 exhibits the accumulation of government assets which occurs in period 1, which is the first period that the government is free to set the capital income tax. In this period, the capital income tax is 214 percent, as shown in Figure 4.

When $T = 2$ the character of the solution has changed rather dramatically. The government is no longer able to lower the value of k_0 by engineering a consumption boom. Nevertheless, it does initially accumulate assets by levying a rather high capital income tax in period T. Hence the frontloading phenomenon remains. As T increases further, the period T capital income tax eventually turns into a subsidy and gives the private sector incentives to invest heavily in the periods leading up to $T-$day ($t = T$). This speeds up convergence of the capital stock to its steady state value, so that at $T = 10$ it is almost immediate. When T increases beyond 15, the solution does not change much when T is increased further.

Similarly, the behavior of consumption and hours as displayed in Figure 2 is best understood in the light of the optimal policies as displayed in Table 3. For example, consider the path for consumption. When $T = 0$, there is a huge boom in period 0, corresponding to the low incentives to save as a result of the very high (214 %) capital income tax in period 1. Subsequently, consumption falls precipitously in response to the fall in the capital income tax rate. Consumption then converges from below to its new, higher steady state as capital is accumulated and consumption possibilities are extended. When $T = 2$, the period 0 consumption boom is gone, and the economy instead commences its accumulation of capital towards the new steady state immediately. In this case, consumption hits its maximum in period 1, the period before the high (97 %) capital income tax rate.

As T increases further, the dip in consumption that corresponds to the capital accumulation is mitigated considerably. This is made possible by the increase in hours worked during the periods $0 \leq t < T$. This increase (as t and T increase) is due to the gap between pre-reform an post-reform labor income tax rates and the accumulation of capital which raises the real pre-tax wage.

6.2 Welfare properties of the solution

As shown in Table 2, the welfare gains diminish as the pre-announcement period T increases. When $T = 0$, the welfare gains correspond to an increase in consumption of about 1.74 percent. When $T = 5$, this figure is 1.16 percent, and when $T = 15$ it has fallen to 0.57 percent. In principle, this is not so strange. Preannouncement (unlike commitment) represents a constraint which a maximizing government would like to avoid if it could. To put the figures into perspective, we recall that CCK find that 80 percent of the welfare gains come from the period 1 capital income tax (and concomitant period 0 consumption boom). They do this by considering a tax reform which sets capital income tax rates to zero in all periods from period 1 on. (This tax reform is of course not optimal.) To compare with the results of the current paper, we note the following. When we perform a similar constant-tax reform experiment to that of CCK, we find that only about 60 (rather than 80) percent of the welfare gains from tax reform disappear. The reason why we get a smaller number than CCK in this case has to do with the fact that their initial stock of privately held assets (capital and government debt) is greater, so that the welfare gains from bringing down its market value are potentially greater.[6]

However, a look at the welfare gains from pre-announced tax reform shows that an even smaller fraction of the welfare gain should be attributed to the high initial capital income tax. When $T = 6$ and the initial capital income tax rate is only about 6 percent, the welfare gains correspond to an increase in consumption of about 1.11 percent, so that only about a third of the welfare gains of the $T = 0$ case have disappeared. When $T = 7$ and the initial capital income tax rate is -6 percent, the welfare gains still correspond to an increase in consumption of about 1.05 percent, so that only about 40 percent of the gains have disappeared. This shows that there is still a rather large free lunch to be had from optimal tax reform even when its

[6] The reason for the initial capital stock in CCK being higher than in the present paper is that our "current" capital income tax rate is higher; we set it to 44 percent, while CCK set it to 27 percent. See section 4.1.1 for the sources of our figures. Meanwhile, we set the initial stock of government debt to zero whereas CCK calibrate it to 51 percent of GDP.

unappealing features have been removed.

This claim receives further support from the following considerations. Our calculations of the welfare gains from optimal pre-announced tax reform presented so far include the effect of mere delay as well as the effect of pre-announcement on the character of the tax reform. Thus it may be that the reduction in welfare gain as a result of pre-announcement is a result of the delay rather than the removal of its unpalatable features. Indeed, as we shall see, that is very nearly the case.

6.2.1 Compensating for the pure delay effect

In this section, we discuss the extent to which the reduction in welfare gain when T increases is a result of mere delay (a good thing coming later rather than sooner), as opposed to the changed character of the optimal tax reform. To investigate this issue, we look again at the allocations generated by the optimal tax reforms for different values of T. Each one of them is then evaluated according to the objective function given by

$$\mathcal{U}_T = \sum_{t=T-7}^{\infty} \beta^{t-T+7} U_t^T,$$

where U_t^T is the period t utility generated by the optimal tax reform announced at $t = 0$ and implemented in period T. Notice that for $T < 7$, the objective function will involve periods with $t < 0$. For these periods, we set the allocations (and hence the period utility) equal to the steady state values generated by current tax policy. Notice that this delay-compensated measure coincides with the usual one when $T = 7$. The choice of the number 7 is motivated by the fact that when $T \gg 0$, the allocation generated by the optimal pre-announced policy stays more or less at the steady state generated by current policy until there are about 7 periods left until the implementation of the reform (see Figures 1 and 2). This means that, in each case, we are evaluating all parts of the path that are affected significantly by the reform. It also means that there are always exactly 7 pre-reform periods in each evaluated path, so that the mere delay effect is held constant.

The results can be seen in Table 4. We notice that once the tax reforms for different values of T are put on a level playing field with respect to the pure delay effect, their welfare gains are rather similar. Thus we conclude that the welfare gains from optimal tax reform do not hinge on initial confiscations or consumption booms. Indeed, the welfare gains are greatest when $T = 7$ when the capital accumulation is smoothest and the initial consumption boom is gone entirely.

7 Extensions

In this section we discuss the robustness of the results to changes in two aspects of the benchmark setup: the period length and the assumption of zero initial government debt.

7.1 Period length

Since there is full commitment, the period length does not represent the frequency with which the government can reconsider its policy. It does, however, represent the frequency with which all variables, including tax rates, can be changed, and it is not obvious that the results are robust to changes in this frequency. Nevertheless, we find that they are. In the benchmark case, the parameters β and δ are calibrated so that the period length can usefully be thought of as 2 years. When these parameters are changed so that each period corresponds to a single year, the results are very similar. As seen in Figures 5 and 6, the dynamics of capital and consumption are largely unchanged except possibly for initial consumption when $T = 0$ which is noticeably higher when the period length is shorter.[7] The difference disappears almost completely when $T = 5$.

7.2 Calibrating initial government debt

In the benchmark parametrization, we set $b_0 = 0$. The results are not affected in any significant way if instead we follow CCK in calibrating initial government debt to equal 51 percent of GDP in the steady state generated by average current policy. Figures 7 and 8 exhibit the dynamics of capital, consumption and hours for $T = 0$ and $T = 5$, with and without positive initial debt. The main qualitative difference lies in the increased incentive to lower the value of initial private assets (since these are now larger) when $T = 0$. This leads to a larger initial consumption boom in this case. When policy is pre-announced 5 periods (10 years) in advance, however, initial government debt barely matters at all for the allocations.

Nevertheless, initial government debt does of course matter for the behavior of government debt itself. As displayed in Figure 9, an increase in initial debt leads essentially to a permanent upward shift in government debt by the same amount,

[7] Notice that equal distances along the horizontal axis in the figures correspond to equal numbers of years rather than equally many time periods.

leaving government deficits almost unchanged.

Finally, it is worth noting that positive initial government debt causes interest payments to crowd out transfer payments. At the given current tax rates and government consumption, transfer payments fall from about 4.6 percent of GDP to approximately zero when initial government debt is increased from zero to 51 percent of GDP.

8 Conclusions

When the government has to pre-announce, and precommit to, its tax policies far in advance, Ramsey optimal reform in the growth model loses some of the rather counter-intuitive properties attributed to it in the literature. In particular, the solution does not feature a consumption boom nor a large capital levy at the time of reform. Moreover, the requirement of pre-announcement can be motivated by using the same arguments as we use to rule out lump-sum taxation.

As a by-product, we have calculated a more reasonable measure of the welfare gains from optimal tax reform by removing the part which can be attributed to an initial policy surprise. We find that only about a third of the welfare gains from optimal tax reform remain when reform is announced 30 years in advance. Yet we also find that the welfare loss of pre-announcement is on the whole a result of the delay itself rather than of the character of the pre-announced reform. When controlling for the pure delay effect, we find that welfare gains of optimal tax reform are roughly the same regardless of the time interval between announcement and implementation.

References

[1] Chamley, C. (1981), "The Welfare Cost of Capital Income Taxation in A Growing Economy," *Journal of Political Economy* **89**, 469-496.

[2] Chari, V.V., L. J. Christiano & P. J. Kehoe (1994), "Optimal Fiscal Policy in a Business Cycle Model," *Journal of Political Economy*, **102**(4), 617-653.

[3] Jones, L. E., R. E. Manuelli and P. E. Rossi (1997), "On the Optimal Taxation of Capital Income," *Journal of Economic Theory,* **73**, 93-117.

[4] Judd, K. L. (1985), "Redistributive Taxation in a Simple Perfect Foresight Model", *Journal of Public Economics*, **28**(1), 59-83.

[5] Lucas, R. E. & N. L. Stokey (1983), "Optimal Fiscal and Monetary Policy in an Economy without Capital", *Journal of Monetary Economics*, **12**(1), 55-93.

[6] Marcet, A. & D. Marshall (1995), "Solving Nonlinear Rational Expectations Models by Parameterized Expectations," Discussion Paper 91, Federal Reserve Bank of Minneapolis.

[7] Marcet, A. & R. Marimon (1995), "Recursive Contracts," manuscript, Universitat Pompeu Fabra and European University Institute.

[8] Ramsey, F. P. (1927), "A Contribution to the Theory of Taxation," Economic Journal, **37**, 47-61.

[9] Turnovsky, S. J. and W. A. Brock (1983), "Time consistency and optimal government policies in perfect foresight equilibrium," *Journal of Public Economics*, **13**, 183-212.

Table 1. Frontloading.

T	Ratio
0	0.09
2	0.05
5	0.03
10	0.02
15	0.02

Footnote: The ratio is 1 minus the steady state ratio between tax revenue and government expenditure.

Table 2. Welfare.

T	Δc
0	1.74 %
2	1.24 %
5	1.16 %
6	1.11 %
7	1.05 %
10	0.85 %
15	0.57 %

Footnote: Δc is the (percentage) increase in consumption in each period required to bring about the same increase in utility as an optimal tax reform.

Table 3. Optimal tax rates.

T	τ^k_T	τ^k_{T+1}	τ^k_∞	τ^h_T	τ^h_{T+1}	τ^h_∞
0	44%	214%	0%	1%	30%	30%
2	97%	1%	0%	31%	31%	31%
5	23%	0%	0%	32%	32%	32%
6	6%	0%	0%	33%	33%	33%
7	-6%	0%	0%	33%	33%	33%
10	-26%	0%	0%	33%	33%	33%
15	-34%	0%	0%	33%	33%	33%

Footnote: When $T = 0$, τ^k_T is constrained to equal 44 %.

Table 4. Welfare gains with the pure delay effect held constant.

T	Δc
0	0.97 %
2	0.81 %
5	0.98 %
6	1.02 %
7	1.05 %

Footnote: Δc is the (percentage) increase in consumption in each period required to bring about the same increase in utility as an optimal tax reform.