European University Institute

MAX WEBER PROGRAMME FOR POSTDOCTORAL STUDIES

EUI WORKING PAPERS

MWP 2014/21
Max Weber Programme

Time-consistent Institutional Design

Charles Brendon and Martin Ellison

European University Institute
Max Weber Programme

Time-consistent Institutional Design

Charles Brendon and Martin Ellison

EUI Working Paper **MWP** 2014/21

ISSN 1830-7728

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu
cadmus.eui.eu

Abstract
We propose a new normative approach to designing institutional commitments in environments that are subject to a time inconsistency problem, à la Kydland and Prescott (1977). This approach captures the idea that institutions should be chosen in a way that is time consistent: if a commitment is found to be best in some sense today, it should remain best in the same sense tomorrow. This property is not satisfied by the usual Ramsey plan, but it can be achieved by placing appropriate restrictions on the choice set of possible commitments. Using a canonical capital tax problem as a laboratory, we consider the implications for institutional design of restricting choice to sets that exhibit this form of time consistency. We show that any optimal plan within a time-consistent choice set must converge to a steady state that differs from the long-run outcome under Ramsey policy. In particular, this outcome exhibits positive long-run capital taxes. This occurs because a time-consistent policy cannot have both high initial capital taxes and zero long-term rates. A policymaker who discounts the future will always be willing to accept long-run distortions in order to tax the inelastic initial capital stock.

Keywords
Institutional design, time inconsistency, commitment policy, capital taxation

JEL codes: E61, E62, H21

This is a substantially revised version of a previous manuscript that circulated under the title 'Optimal Policy behind a Veil of Ignorance'. We thank Árpád Ábrahám, Paul Beaudry, Tatiana Damjanovic, Paul Levine, Albert Marcet, Ramon Marimon, Alex Mennuni, Meg Meyer, Michel Normandin, Ctirad Slavik, Rick Van der Ploeg, and Simon Wren-Lewis for detailed discussions and comments, together with numerous seminar participants. All errors are ours.

Charles Brendon
Max Weber Fellow, 2012-2013 and 2013-2014
European University Institute. Email: charles.brendon@eui.eu

Martin Ellison
University of Oxford. Email: martin.ellison@economics.ox.ac.uk

1 Introduction

Time inconsistency problems are near-ubiquitous in macroeconomic policy design. Whenever agents' expectations of the future influence what a government can do today, and whenever these expectations can in turn be affected by future policy choices, the actions of policymakers in the future act as a constraint on the present. Unless future policymakers are institutionally constrained by accumulated obligations, they will have no reason to consider the effects of their choices on past expectations. The resulting outcome need not lie on the dynamic Pareto frontier: policymakers in every time period could be made better off under an alternative feasible plan. This troubling feature of discretionary choice was, of course, first formalised by Kydland and Prescott (1977), who inferred from it the broad recommendation that repeated choice was to be avoided where possible. Policy should be based on 'rules rather than discretion'.

The main focus of the current paper is on how these rules – 'institutions', as we label them – should be designed. This focus is unusual, as the main distinction drawn in the literature is between governments that can commit and governments that cannot.[1] Once the fact of commitment is given, the assumption is that it will be to a Ramsey-optimal plan – i.e., the best feasible allocation from the perspective of the first period of the model.[2] We instead define and characterise the new concept of 'time-consistent institutional design', and show that it has a number of features that make it an appealing alternative to Ramsey policymaking.

Ramsey plans have two related features that make them unlikely practical solutions to institutional design problems. First, they are time-contingent. Even in purely deterministic models, the Ramsey-optimal policy choice for a given state of the economy will vary depending on the number of time periods that have elapsed since optimisation first took place. The best capital tax rate for period 0 is different from the best capital tax rate for period $t > 0$, even if the capital stock in these two periods is identical. This means that Ramsey policy cannot be viewed as a set of 'rules' in the common sense of the term: that is, time-invariant functions mapping from the observed state of the economy to an optimal choice. As a consequence, apparently straightforward economic policy questions – such as: 'What is the best rate of inflation to target?', or: 'How should taxes respond to the economic cycle?' – are not answered directly by the solution to a Ramsey plan. At the same time, Markovian discretionary policy – of the sort recently analysed by Klein et al. (2008), among others – takes as

[1] The literature contrasting commitment and discretionary outcomes in different settings is huge, and we do not attempt to do it justice here. Recent contributions include – but are not limited to – Klein, Krusell and Ríos-Rull (2008), Debortoli and Nunes (2013), Song, Storesletten and Zilibotti (2012), Yared (2010), Martin (2009), Díaz-Giménez et al. (2008) and Ellison and Rankin (2007).

[2] An intermediate point is found by the literature on 'loose commitment', in which the policymaker is permitted to re-optimise in each period with a given probability: see, for instance, Debortoli and Nunes (2010) and Schaumburg and Tambalotti (2007). Nonetheless, this literature retains the perspective that when commitment plans are designed, they are designed to be best from the perspective of the current time period.

given that institutional commitment is impossible. It need not have particularly desirable welfare policies, for the well-known reasons first emphasised by Kydland and Prescott. Given the unsatisfactory character of both of these options, the danger is that important policy questions may consequently be resolved in an *ad hoc* fashion, without clear reference to a meaningful optimality concept – even in environments where well-established economic models would be ideally suited to giving guidance. Our main motivation for developing the idea of time-consistent institutional design is precisely to offer a more realistic framework for optimal policy choice, allowing good economic models to have greater practical influence. As discussed in Section 2, we are not the first authors to seek 'justified pragmatism' of this kind, but to date there remains no generally-accepted approach.

To be clear, the time-contingency property just discussed is a feature of Ramsey policy notwithstanding the fact that a Ramsey plan may be represented as solving a recursive choice problem. The works of Marcet and Marimon (2014) and Abreu, Pearce and Stachetti (1990) in particular have greatly expanded the analytical toolkit available when solving for Ramsey policy – showing how a state vector augmented with either Lagrange multipliers or promise values delivers recursivity. But in each case the value of this augmented state vector in a given time period will depend on the time that has elapsed since optimisation. The first period of the problem, for instance, will never feature binding obligations related to past promises, whereas later periods will. The natural state vector alone will thus not be enough to determine policy choice in a given period.

The second, closely-related feature that makes Ramsey policy unappealing from a normative perspective is that it is time-inconsistent. An institutional designer fixing Ramsey policy from period 0 onwards will choose a different set of outcomes for time period $t > 0$ than an institutional designer fixing Ramsey policy from a later period $s \leq t$ onwards. This is troubling when the Ramsey perspective is used to analyse public policy – rather than, say, dynamic contracts among private agents – because legislative bodies in many countries simply do not have the right to bind their successors.[3] If a subsequent government were to re-assess a previous Ramsey commitment according to the same Ramsey paradigm, it would find that the previous commitment is no longer best. It is hard to see how an institution of this kind could come into being in some initial 'period 0', but face no subsequent threat to its continued existence. As Lars Svensson (1999) has put it: "Why is period 0 special?".

Like time contingency, time inconsistency remains present despite the fact that the Ramsey choice problem can be made recursive via an expanded state vector. *Ex-post*, policymakers will always have an incentive to 'reset' the elements of the state vector that correspond to promises – fixing the inherited vector of multipliers to zero, or the inherited promise values to a level such that any associated constraints are slack. Again, in a microeconomic dynamic contracting environment this will not be a problem so long as contracts can be

[3] This is especially true of countries such as the United Kingdom in which the legislature is legally sovereign, and not subject to oversight by any constitutional court.

enforced: legal punishments can effectively make inherited promises 'natural' states. But dynamic contracts are not the relevant institutional framework when thinking about public policy.

The time inconsistency of Ramsey plans is an unavoidable feature of Kydland and Prescott problems. But at the same time, the Ramsey framework allows for very substantial freedom of choice on the part of the institutional designer, and in this regard it may be posing an unrealistic problem. For every time period, and – where relevant – every history of exogenous shocks, a different promise can be chosen. This is very different from the problem of, say, choosing a single inflation target to endure for a large number of years. In that case there is a straightforward cross-restriction on the promises that can be set in different periods: they must all equal one another. Such cross-restrictions are arguably the norm rather than the exception when long-term economic policy is being legislated. It is extremely rare in practice for directly time-contingent policies to be committed to *ex ante*.

The main point of our paper is that time inconsistency is directly connected to the freedom of choice given to a Ramsey planner. We show how appropriately-chosen cross-restrictions on the set of possible dynamic commitments can result in an institutional design problem that – unlike the Ramsey plan – is time-consistent. That is, the optimal commitment from this restricted set of institutions in period 0 would still be judged optimal, within the same restricted set, if it were reassessed at a later date.[4] This focus on restricted choice sets can be rationalised normatively as a requirement that no generation should be granted special privileges relative to its successors. The Ramsey plan can always be viewed as a solution that assumes the period-0 institutional designer is granted unrestricted choice, but this plan would survive *ex-post* reassessment only if later generations faced greatly restricted choice sets by comparison. We are interested in the consequences of *symmetric* restrictions on the choice sets that institutional designers of different generations are allowed, whilst still allowing as much meaningful scope for choice as possible.

We study the properties of institutions that are optimal time-consistent choices of this sort, using as our laboratory a deterministic capital tax problem in the spirit of Chamley (1986) and Judd (1985), with a simplifying assumption that the government must run a balanced budget. This is an ideal environment, as it is a model in which Ramsey policy involves particularly acute asymmetries over time: initial-period capital taxes are always implausibly high, but in the long run they converge to zero.

In this context we associate institutional choice more specifically with choice over dynamic sequences of promise values. These promise values are then treated as exogenous restrictions on a 'day-to-day' policy choice problem. The promise values are real numbers, corresponding to minimal levels of post-tax wealth that

[4] The idea of 'reassessment' here differs from actual choice, as we do not want to reduce the institutional design problem to a Stackelberg game. We assume that institutional designers always assess options under the assumption that their chosen institution will endure indefinitely. The question is whether the best choice under this assumption will be a time-consistent one.

the consumer must be allowed to possess each period in a market equilibrium. When they constrain day-to-day choice, it is to rule out arbitrarily high capital taxation. Ramsey policy follows from choosing these promises optimally from the unrestricted choice set $\mathbb{R}^{\infty}$, under the period-0 welfare criterion. Our perspective is still that they should be chosen optimally from the perspective of period 0, but that this selection should instead be from a restricted choice set within which choice is time-consistent. This is what we call time-consistent institutional design.

1.1 Preview of main results

The requirement for time-consistent institutional choice restricts the problem considerably, but not completely. In particular, our first main result is that a time-consistent institutional design problem can have at most one meaningful degree of freedom for policy choice. Thus the cross-restrictions that link promises in different periods must have the effect of reducing the institutional designer's choice set to – at most – a family of elements of $\mathbb{R}^{\infty}$, indexed by just one parameter. Choice over this parameter is the only decision that a time-consistent institutional designer faces. This is obviously a very restricted choice by comparison with the Ramsey planner's, but it is not wholly unrealistic. It implies the same degree of freedom that policymakers accept willingly whenever they legislate for policy variables that are not time-varying.

Having demonstrated that time-consistent choice is at most unidimensional, we next begin to place structure on this choice. We make three parsimonious assumptions relating to the institutional designer's choice set, capturing the idea that this set should allow as much scope for choice as possible. Specifically, we assume that the choice set contains one degree of freedom everywhere, and that it allows for meaningful variation in promises at all horizons as the single policy parameter is varied. We show that these assumptions together are enough to pin down a unique steady-state outcome for time-consistent policy in the given optimal tax model. Importantly the capital tax rate associated with this steady state will always be positive, irrespective of parameter values. It is 21 per cent for our benchmark case. This is in sharp contrast with the steady-state of Ramsey-optimal policy, which sets capital taxes to zero. The reason for the difference is that when choice is limited to one degree of freedom, the policymaker must trade off the relative benefits of high initial capital taxes and low steady-state rates. Even though capital taxes may adjust over time as the state of the economy evolves, positive initial tax rates cannot be selected in a time-consistent choice set unless they are accompanied by positive steady-state rates. The result that the steady state of Ramsey policy does not coincide with the steady state from optimal time-consistent choice is very general, and calls into question the intrinsic desirability of the Ramsey steady state – viewed independently of the transition to it.

Finally we consider time-consistent capital taxes on the transition path as the capital stock accumulates. For this we need to place significantly more structure on the choice problem: the parsimonious assumptions that determine steady state are too weak to fix a unique transition. We add two further restrictions – the first ensuring that for a given capital stock the chosen policy will always be the same, irrespective of time, and the second that no policymaker 'enivies' another, in the sense of preferring the optimal promise sequence available when the capital stock takes a different value. When these restrictions hold, we show that there is a single possible time-consistent institutional choice, and that it sets the promise value equal to a constant in all time periods. This is not as restrictive as it first appears, for reasons that we explain. The exact character of capital and labour taxes along the resulting transition will depend on the precise parameterisation, but in all cases the departures of taxes from their steady-state values are an order of magnitude smaller than under Ramsey policy. And unlike Ramsey policy, if the initial capital stock is at its steady-state level, all variables in the economy remain constant indefinitiely.

The rest of the paper is structured as follows. Section 2 places our contribution in the broader policy design literature. Section 3 presents a simple deterministic social insurance problem without state variables, which we use to motivate our broader search for time-consistent institutions. Section 4 introduces the main capital tax problem that we study, and decomposes it into separate 'day-to-day' and 'institutional design' problems. Section 5 then formalises our notion of time-consistent institutional design, and characterises its implications for steady-state outcomes. Section 6 considers transition dynamics, and section 7 concludes. An appendix collects proofs.

2 Related literature

Our concept of time-consistent institutional design is – to the best of our knowledge – novel, but a number of strands of the policy design literature relate closely to what we do. A different concept of time-consistent public policy has received significant recent attention in the literature. This focuses on the positive analysis of Markovian equilibria in the dynamic Stackelberg game among successive policymakers when commitment is not possible. Recent papers to apply this perspective in various settings include Klein and Ríos-Rull (2003), Ortigueira (2006), Klein, Krusell and Ríos-Rull (2008), Diaz-Gimenez et al. (2008), Martin (2009), Blake and Kirsanova (2012), Reis (2013) and Niemann, Pichler and Sorger (2013). The basic assumption of this literature is that policy is chosen sequentially as a function of the economy's natural state vector, with every generation of policymakers treating the response function of its successors as given. We instead suppose that policy is chosen once-and-for-all at the start of time, but from a restricted set that is constructed to ensure that this choice would still be considered optimal if it were reassessed at a later date. Our notion of time consistency therefore relates to the once-and-for-all institutional design choice, not sequential choice on the part of successive discoordinated generations.

Like this literature we obtain policies that are functions only of the aggregate state of the economy, at least in deterministic settings of the type considered in this paper.[5] But unlike it, our work is explicitly normative: the dynamic allocations that we generate can meaningfully be described as *optimal* time-consistent plans, albeit subject to the constraints required for time consistency. They are the outcomes of constrained choice over entire dynamic policy sequences, not of isolated, strategic period-by-period choice.

The concept in the literature that comes closest to our notion of time-consistent institutional design for Kydland and Prescott problems is the 'timeless perspective' of Woodford (1999, 2003, 2010). Like ours, Woodford's approach aims to address the fact that practical policy design does not generally consider time-contingent policy of the sort that is Ramsey-optimal. Since Ramsey-optimal policy will generally converge to a steady state, his proposal is to implement this steady-state allocation from the start of time. The normative justification for doing so is that such a policy would have been best for a policymaker optimising in the distant past.

Woodford's approach has initiated a useful discussion on how best to design time-non-contingent policy commitments, to which this paper is a contribution. But unfortunately it has a number of weaknesses. In general it is not clear why the long-run outcome of Ramsey policy should be viewed as desirable, independently of the transition path to steady state. In the simplest possible time inconsistency problems – deterministic settings with no state variables – the timeless perspective policy will be completely constant over time, but it will be welfare-dominated in the set of such constant policies. This point is made more fully in section 3 below. It relates to a similar point made by Blake (2001) in the context of a stochastic New Keynesian stabilisation problem, where the timeless perspective policy was shown to generate a higher expected value for the loss function than a specific alternative. In this paper we restrict our focus to deterministic environments. If there are additionally no state variables, the optimal time-consistent institution that we study will deliver the best constant policy, unlike the timeless perspective.

A similar – though less formal – focus on Ramsey steady-state outcomes has also characterised much of the literature on optimal capital taxation in complete markets. Atkeson, Chari and Kehoe (1999), for instance, focus their policy advice exclusively on the long-run result of a zero tax rate in Chamley-Judd settings, arguing informally that lags in implementation would render optimal transition dynamics irrelevant: the steady-state of optimal policy should be targeted immediately. We instead find that a positive steady-state tax should be chosen by a time-consistent institutional designer. A similar finding has previously been obtained either by relaxing the complete markets assumption of the standard Ramsey approach,[6] or by conducting the analysis in an overlapping generations setting,[7] or both.[8] Our analysis instead retains the complete-markets structure of the classic Chamley-Judd setting, with a single infinitely-

[5] In stochastic models it will be desirable to condition on the history of exogenous shocks. We neglect this dimension in the present work for simplicity.

[6] For instance, Aiyagari (1995) and İmrohoroğlu (1998).

[7] For instance, Alvarez et al. (1992) and Erosa and Gervais (2002).

[8] See, in particular, Conesa, Kitao and Krueger (2009).

lived representative agent. The reason we obtain positive steady-state capital taxes is that time-consistent policy forces a trade-off between initial capital tax rates and their long-run levels. It is only by accepting long-run distortions that a time-consistent institutional designer can tax away inelastic capital in period 0. This trade-off is not captured appropriately by policy advice that recommends an immediate jump to the Ramsey steady state.

3 An example without states

Before considering the optimal capital tax model that will be our main focus, it helps to fix some ideas in a simple setting without state variables. Consider a two-period deterministic overlapping generations economy in the style of Samuelson (1958). Within each generation, half of the agents are born lucky and receive a relatively high lifetime endowment, whilst half are born unlucky and receive a relatively low lifetime endowment. These endowments, which are publicly observable, are denoted $\left\{\overline{y}^{1,h}, \overline{y}^{2,h}\right\}$ for the lucky agent and $\left\{\overline{y}^{1,l}, \overline{y}^{2,l}\right\}$ for the unlucky agent:.[9] These endowments are assumed to satisfy the four inequalities: $\overline{y}^{1,h} > \overline{y}^{2,h}$, $\overline{y}^{1,l} > \overline{y}^{2,l}$, $\overline{y}^{1,h} > \overline{y}^{1,l}$ and $\overline{y}^{2,h} \geq \overline{y}^{2,l}$. Thus both agents receive higher endowments when they are young, which allows for the well-known possibility that market outcomes may be dynamically Pareto inefficient.

Individual lifetime utility functions are common for all agents in all generations, and take a standard time-separable form:

$$U\left(c^1, c^2\right) = u\left(c^1\right) + \beta u\left(c^2\right) \tag{1}$$

where c^1 is the agent's consumption when young and c^2 when old, and $u(\cdot)$ is increasing and strictly concave. A paternalistic government exists that has preferences over the welfare of all agents in all current and future time periods. These preferences – assessed in some period $t \geq 0$ – are described by the objective W_t:

$$W_t := \sum_{s=t}^{\infty} \beta^{s-t} \sum_{i \in \{h,l\}} \left[u\left(c_s^{1,i}\right) + u\left(c_s^{2,i}\right)\right] \tag{2}$$

where $c_s^{1,h}$ is the consumption of the young, high-endowment agent in period s, and so on.

The government wishes to commit to a permanent social security scheme that will transfer resources from rich to poor, subject to the restriction that participation in the scheme will be voluntary – so in each period all agents must be left at least as well off as they would be if they were to opt out, and consume their endowment alone. This implies participation constraints for all types in all periods. We will focus on the case in which the participation constraint of

[9] Superscripts 1 and 2 denote the relevant time period of the agent's life, and h and l denote high and low values respectively.

young, high-endowment agents is the only binding restriction. This constraint, in generic period t, is:

$$u\left(c_t^{1,h}\right) + \beta u\left(c_{t+1}^{2,h}\right) \geq u\left(\overline{y}^{1,h}\right) + \beta u\left(\overline{y}^{2,h}\right) \tag{3}$$

In addition to this restriction, there is an aggregate resource constraint that must hold in each time period:

$$c_t^{1,h} + c_t^{2,h} + c_t^{1,l} + c_t^{2,l} \leq \overline{y}^{1,h} + \overline{y}^{2,h} + \overline{y}^{1,l} + \overline{y}^{2,l} \tag{4}$$

3.1 Institutional design, and day-to-day policy

Despite the relative simplicity of the setup, for the purposes of the subsequent discussion it is instructive to divide the participation constraint (3) into two parts: a 'promise-making' restriction and a 'promise-keeping' restriction. This is just a matter of defining a sequence $\{\omega_t\}_{t=0}^{\infty}$ such that for all $t \geq 0$ we have:

$$u\left(c_t^{1,h}\right) + \beta\omega_{t+1} \geq u\left(\overline{y}^{1,h}\right) + \beta u\left(\overline{y}^{2,h}\right) \tag{5}$$

and:

$$u\left(c_t^{2,h}\right) \geq \omega_t \tag{6}$$

where ω_{t+1} is thus the amount of utility that has been promised to old, high-endowment agents in period $t+1$. The version of constraint (6) for period 0 is not a fundamental restriction on optimal commitment strategies from period 0 onwards, and when selecting Ramsey policy we can always let ω_0 take a sufficiently low value that it is not binding. We nonetheless include the period-0 restriction in our description of the constraint set, as our concept of time-consistent institutional design will admit the possibility that it could bind. For notational ease we hereafter use the notation $\omega^{0,s}$ to refer to a finite sequence of promises from period 0 to period s, $\{\omega_t\}_{t=0}^{s}$, and $\omega^{0,\infty}$ to denote an infinite sequence $\{\omega_t\}_{t=0}^{\infty}$.

We consider that a policymaker has access to a commitment device when they are able to fix an infinite sequence $\omega^{0,\infty}$, such that the promise-keeping constraint is binding in some or all periods. Our main focus will be on alternative approaches to determining such sequences.

For any given $\omega^{0,\infty} \in \mathbb{R}^{\infty}$, we can consider the problem of how best to choose a consumption allocation such that these promises are satisfied at all horizons. In the initial period 0 this takes the form:

$$\max_{\left\{c_t^{1,l}, c_t^{1,h}, c_t^{2,l}, c_t^{2,h}\right\}_{t=0}^{\infty}} \sum_{s=0}^{\infty} \beta^s \sum_{i \in \{h,l\}} \left[u\left(c_s^{1,i}\right) + u\left(c_s^{2,i}\right)\right]$$

subject to the resource constraint (4), the promise-making constraint (5), the promise-keeping constraint (6), and $\omega^{0,\infty}$. The constraint set for this problem is easily seen to be convex, and the objective criterion is concave. Provided feasibility can be satisfied, it will thus have a unique solution. This problem can be thought of as one solved by a 'day-to-day' bureaucratic policymaker, administering the practical aspects of choice under a given institution – represented by the fixed sequence of promises. We denote the value of the allocation that solves this problem $V_0^d\left(\omega^{0,\infty}\right)$.[10] Importantly, the allocation that solves the day-to-day problem is entirely time-consistent. Inconsistency in general derives from the incentive to reset promise values that have previously been fixed, which – by construction – does not apply to a choice problem that abstracts from the determination of the ω_t terms.

Formally, time inconsistency can be viewed as a change over time in the policymaker's preference ranking over promise sequences. The function $V_0^d\left(\omega^{0,\infty}\right)$ describes a complete, reflexive and transitive binary preference relation over the space of all $\omega^{0,\infty} \in \mathbb{R}^\infty$, labelled $\succsim_0$. The same function can be used to describe a conditional preference ordering over promises from t onwards, $\omega^{t,\infty}$, given some assumption about the determination of $\omega^{0,t-1}$. We will focus on the case in which, for the given choice of $\omega^{t,\infty}$:

$$\omega^{0,t-1} \in \arg \max_{\left\{\omega^{0,t-1}:\left\{\omega^{0,t-1},\omega^{t,\infty}\right\}\in\Omega_0\right\}} V_0^d\left(\omega^{0,t-1},\omega^{t,\infty}\right)$$

with $\Omega_0 \subseteq \mathbb{R}^\infty$ a constraint set from which promise sequences must be chosen, explained fully below. Generically these preferences are labelled $\succsim_{0,t}$. Likewise, a relation $\succsim_1$ will describe the preferences in period 1 over commitments from period 1 onwards, mapped by $V_1^d\left(\omega^{1,\infty}\right)$, with $\succsim_{1,t}$ taking a corresponding conditional definition, and so on.

The essence of the time inconsistency problem is that these preference rankings do not coincide. In particular, the preference ranking $\succsim_{0,t}$ on $\mathbb{R}^\infty$ will not coincide with $\succsim_t$:[11] desirable plans from t onwards viewed from the perspective of period 0 need no longer be desirable when viewed from the perspective of period t. Notably, the policymaker in period t will not see the value in allowing the promise-keeping constraint for period t to bind, whereas the policymaker in period 0 will benefit from faithfulness.

The institutional design problem is how best to select a sequence $\omega^{0,\infty} \in \mathbb{R}^\infty$, given time inconsistency. As in the Ramsey paradigm, we will assume that this problem is solved 'once and for all' in period 0, based on the period-0 the preference relation $\succsim_0$. Importantly, however, this choice will not be unrestricted. Analytically the institutional design problem we study takes the form:

$$\max_{\omega^{0,\infty}\in\Omega_0} V_0^d\left(\omega^{0,\infty}\right) \tag{7}$$

where $\Omega_0 \subseteq \mathbb{R}^\infty$ is a constraint set from which the promises can be selected.[12] Alternative structures for this constraint set will be the main focus of our analysis.

[10] If the constraint set is empty we set $V_0^d\left(\{\omega_t\}_{t=0}^\infty\right) = -\infty$.

[11] That is, there will be ω' and ω'' in $\mathbb{R}^\infty$ such that $\omega' \succsim_{0,t} \omega''$ but not $\omega' \succsim_t \omega''$.

[12] The 0 subscript denotes the period in which institutional choice is being made. This will be of relevance when we consider time consistency.

The separation between institutional design and day-to-day policymaking is made more for heuristic than mathematical reasons. Combining the two problems into a single, joint choice over real allocations and promises would deliver the same policy, since the separated formulation takes the 'max-max' form. But since our focus is on the idea of institutional design *per se*, it is useful to be able to isolate a problem that captures exclusively this.

3.2 Optimal choice

It is straightforward to characterise optimal choice in this setting, both analytically and numerically. Consistent with the distinction that we have just drawn, we first specify optimal choice for the 'inner' day-to-day problem, given an arbitrary $\omega^{0,\infty}$ sequence. We then consider different approaches to solving the 'outer' institutional design problem.

3.2.1 Day-to-day choice

Given $\omega^{0,\infty}$, and placing multipliers $\beta^t \mu_t$, $\beta^t \lambda_t^1$ and $\beta^t \lambda_t^2$ on constraints (4), (5) and (6) in turn, the optimal day-to-day policy choice is characterised by:

$$\mu_t = u'\left(c_t^{1,h}\right)\left(1+\lambda_t^1\right) \tag{8}$$

$$\mu_t = u'\left(c_t^{2,h}\right)\left(1+\lambda_t^2\right) \tag{9}$$

$$\mu_t = u'\left(c_t^{1,l}\right) \tag{10}$$

$$\mu_t = u'\left(c_t^{2,l}\right) \tag{11}$$

together with complementary slackness. These are the standard requirements for optimal cross-sectional resource allocation: the weighted value of the marginal utility of consumption is equalised across all agents, with the effective Pareto weight of high-type agents being increased above the underlying utilitarian value of 1 whenever constraints (5) and (6) bind.[13]

3.2.2 Institutional design: Ramsey policy

We now turn to the institutional design problem. Under Ramsey policy there is no restriction on the set of promise values from which choice can be made: $\Omega_0 = \mathbb{R}^\infty$. A standard application of the envelope theorem yields the following necessary optimality condition for choice of ω_t:

$$\lambda_t^1 = \lambda_{t+1}^2 \tag{12}$$

for all $t \geq 0$ and:

$$\lambda_0^2 = 0 \tag{13}$$

[13] Marcet and Marimon (2014) and Mele (2011) have emphasised this interpretation of shadow values on incentive constraints as contributing to Pareto weights.

These requirements are very familiar from the work of Marcet and Marimon (2014) and others. Two aspects are worth emphasising. First, the policymaker finds it optimal to let any increase in the Pareto weight of high-endowment agents be fully persistent over time, so that it remains at its initial value, λ_t^1, even when the participation constraint has ceased to bind. This persistence property follows from the optimal spreading of resources over time, and is a very common feature of Ramsey-optimal policy in the class of models with binding participation constraints. Second, as mentioned already, the Ramsey policymaker never has any incentive to let the initial promise ω_0 constrain choice, since old, high-type agents in period 0 are not at risk of exiting the scheme. From (9) this means that the consumption level of old, high-type agents in period 0 will equal that of low-type agents, and this will relax somewhat the initial resource constraint.

These features are evident in the dynamic paths of consumption for the different agents under the Ramsey plan, which we plot in Figure 1 for an illustrative calibration.[14]

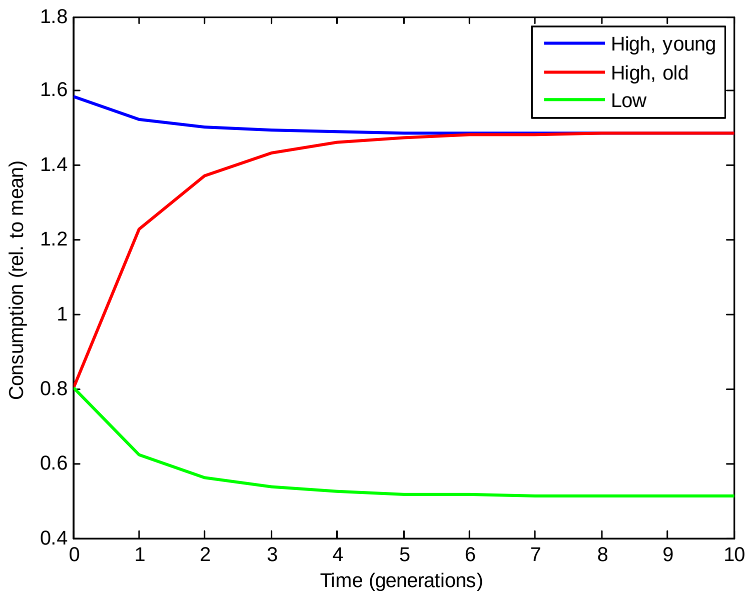

Figure 1: Ramsey-optimal consumption dynamics

The Figure confirms the general lesson discussed in the introduction: despite

[14]Specifically, we assume $\overline{y}^{1,h} = 2.5$, $\overline{y}^{2,h} = 0.625$, $\overline{y}^{1,l} = 0.625$ and $\overline{y}^{2,l} = 0.25$, giving a mean endowment of 1. Consumption utility is assumed to be CES, with elasticity of intertemporal substitution set to 0.5. We assume $\beta = 0.96^{30} \simeq 0.29$, reflecting a four-per-cent annual interest rate and a 30-year gap between generations.

the fact that the economic environment is completely stationary, the Ramsey-optimal social security scheme features non-trivial dynamics. The relatively slack resource constraint in period 0 means that it is optimal initially to front-load the incentives provided to young, high-endowment agents. This imparts further dynamics to the optimal plan, as old, high-endowment types in period 1 will have received a very high allocation in period 0, and so will not require so great a transfer in period 1. This in turn frees up resources in period 1, allowing the incentives of the new generation of high-endowment agents again to be frontloaded. This dynamic continues, and steady state is reached only gradually. In steady state all high-endowment agents, young and old, receive the same high consumption level, and all low-endowment agents receive the same low consumption level.

3.2.3 Institutional design: constant policy

As discussed in the introduction, in many environments the time-contingent character of Ramsey policy makes it an impractical choice for actual delegation policies. This motivates a search for more realistic alternatives – a project initiated by Woodford (1999). The useful feature of the current setting, which is both deterministic and free from natural state variables, is that once intrinsic dynamics in the promises have been eliminated the chosen policy must be entirely constant. There are no shocks or wealth accumulation dynamics that might provide an alternative source of variation. The characteristics of optimal constant policy are very easy to analyse, and provide an unambiguous basis from which to generalise to more complex problems.

Suppose, therefore, that the choice set for the institutional design problem were to include only dynamic promise sequences that are constant. That is:

$$\Omega_0 = \left\{ \omega^{0,\infty} \in \mathbb{R}^\infty : \omega_s = \omega_t, \forall s, t \geq 0 \right\}$$

What are the implications of maximising $V_0^d \left(\omega^{0,\infty} \right)$ on this restricted set? Defining multipliers as before, the following necessary condition is easily obtained:

$$\sum_{t=0}^{\infty} \beta^{t+1} \lambda_t^1 = \sum_{t=0}^{\infty} \beta^t \lambda_t^2 \tag{14}$$

With constant promises and a stationary environment, the outcome of the day-to-day problem will clearly be a set of four constant consumption values, $\left\{ \overline{c}^{1,h}, \overline{c}^{2,h}, \overline{c}^{1,l}, \overline{c}^{2,l} \right\}$, and it follows from this that the multipliers will also be time-invariant. Condition (14) then reduces to:

$$\beta \overline{\lambda}^1 = \overline{\lambda}^2 \tag{15}$$

This implies a subtle difference in the treatment of past obligations between a Ramsey-optimal institution and an optimal constant institution. The effective Pareto weight of old, high-endowment agents is now lower than for young, high-endowment agents: $\left(1 + \beta \overline{\lambda}^1\right)$ as opposed to $\left(1 + \overline{\lambda}^1\right)$. Under the Ramsey plan these weights were constant for the two periods of the agent's lifetime – though potentially differing across generations – which in turn ensured identical steady-state consumption levels for young and old high-endowment agents. Now steady state is imposed from the start of time, but the allocations differ.[15] The promise values associated with the steady state of Ramsey-optimal policy are clearly available to the designer of the optimal constant institution, so it follows that the Ramsey steady state is welfare-dominated in the set of time-invariant policies. This is an important result, despite the simplicity of the environment: it strongly suggests that the steady state of Ramsey policy may not be an appropriate starting point when designing more realistic commitments than the Ramsey plan.

To illustrate the differences, Figure 2 superimposes the optimal constant allocation on the Ramsey-optimal dynamics. Under the constant institution, the consumption of young, high-endowment agents is now around 60 per cent greater than that of old, high-endowment agents. The consumption of low-endowment agents is higher under the optimal constant institution than the Ramsey-optimal institution in every period except for the first.

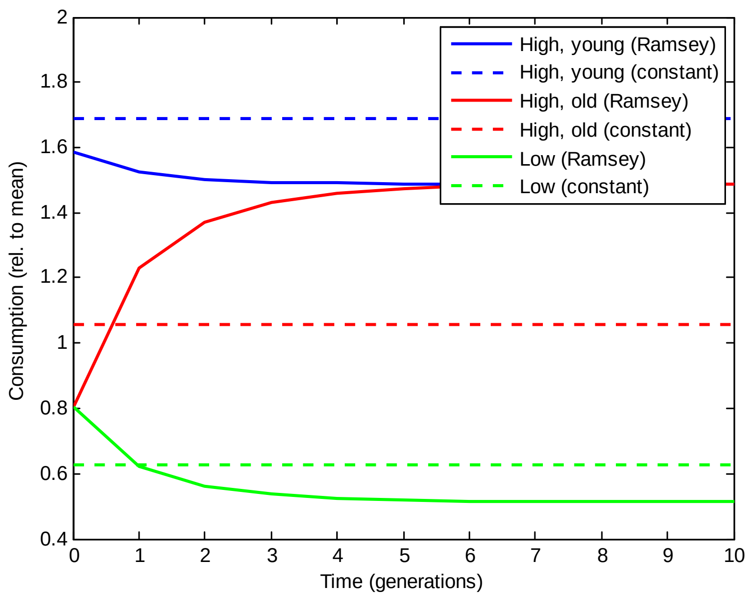

Figure 2: Optimal constant consumption dynamics

Assessed in terms of the social welfare criterion W_0, the Ramsey-optimal allocation obviously dominates the constant policy: by definition it is the best

[15] They would coincide only if $\beta \rightarrow 1$.

allocation when $\omega^{0,\infty}$ is completely unrestricted. But if one compares the continuation value of the two policies over time, the constant allocation is superior after just one (generation-length) period. Figure 3 demonstrates this. It charts over time the permanent, uniform consumption level that is equivalent in social welfare terms to the continuation value of each of the two policies.[16].

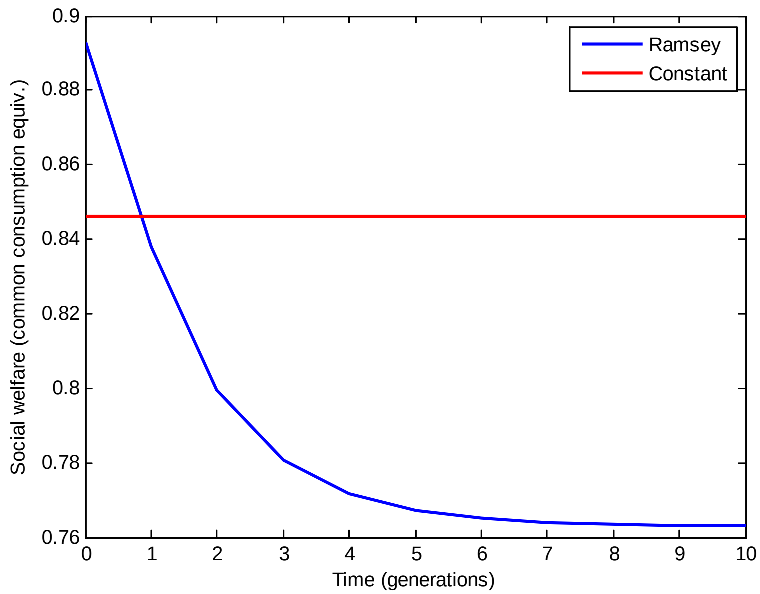

Figure 3: Welfare comparison: Ramsey vs optimal constant policy

3.3 The time-consistency of constant choice

For the purposes of this paper there are two main points that this simple example highlights. The first is that a Ramsey steady state is only desirable as part of a complete dynamic Ramsey plan. It need not have particularly favourable welfare properties *per se*, and this is certainly enough to question its relevance as a focal point in more complex environments with state variables and stochastic shocks.[17] The 'timeless perspective' of Woodford (1999) would recommend implementing this Ramsey steady-state from period 0 onwards, but this is clearly welfare-dominated in the set of constant policies. This motivates searching for

[16]So for a given period t and a given dynamic allocation from t onwards, Figure 3 shows the consumption level (relative to the average endowment value) such that if all agents of all generations were to receive this consumption value in all periods from t onwards, the continuation value of social welfare would be the same as under the chosen policy.

[17]For the remainder of the paper we will focus on an example with the former – leaving a full consideration of policy in stochastic environments to further work.

an alternative perspective on time-non-contingent choice that – like the timeless perspective – will be applicable in more complex environments, but does not suffer from so clear a normative failure.

The second important feature of the example – central to the arguments that follow – is the time-consistency of the constant institutional design problem. Since there are no state variables responding to institutional choice, the best constant promise to enact from period 0 onwards must be the same as from any arbitrary period t onwards. Indeed, the entire preference ordering over constant promises is unchanging through time. Formally, let Ω_t be the set of promises permitted from t onwards under Ω_0:

$$\Omega_t = \left\{\omega^{t,\infty} \in \mathbb{R}^\infty : \left(\omega^{0,t-1}, \omega^{t,\infty}\right) \in \Omega_0 \text{ for some } \omega^{0,t-1} \in \mathbb{R}^t\right\}$$

If Ω_0 is the set of constant promises, clearly so too is Ω_t. Now consider the conditional preference ranking $\succsim_{0,t}$ on Ω_t. Given the stationarity of the problem, it is immediate that this must coincide with $\succsim_t$. That is, rankings in Ω_t from the perspective of an institutional designer in period 0 are the same as from the perspective of an institutional designer in period t.

Importantly, Ramsey policy does not satisfy this form of time consistency. Suppose $\Omega_0 = \mathbb{R}^\infty$. It is clear from Figure 2 that if a policymaker were to select an optimal Ramsey plan from period $t > 0$ onwards, allocations would differ relative to continuing with the optimal period-0 plan. Hence the ranking $\succsim_{0,t}$ over $\Omega_t = \mathbb{R}^\infty$ must differ from $\succsim_t$, and time consistency must be violated. In order to justify a Ramsey plan as an optimal institution, therefore, one must be willing to accept that the policymaker in period 0 has a unique privilege relative to his or her successors: the freedom to impose a choice of promises that will be regretted *ex post*, given the same set of options. When time consistency obtains, there are no such regrets.

This notion of time consistency as 'lack of regret' differs from the notion commonly associated with discretionary policymaking, and studied by Klein et al. (2008), Díaz-Giménez et al. (2008) and others. In these papers the assumption is that institutional commitments of any form are impossible, meaning that promise-keeping constraints can never bind. Private-sector expectations fully account for this, and 'time-consistent policy' in this literature is the outcome of the Stackelberg game that results. We instead allow for a perfect commitment device, so that expectations can be fully influenced by institutional choice, and past promises can bind. We seek time consistency in the choice over complete sequences of promises, not of period-by-period allocations.

There is an undeniable tension in the requirement that a *commitment* strategy should be time-consistent in the sense we describe. The whole point of commitment is that it should provide a 'cast-iron' policy plan, unchallengeable by future generations regardless of their preferences. If this is the case, why bother requiring that choice should not be regretted? Our emphasis here again is pragmatic. In practice governments cannot tie the hands of their future selves perfectly, though they can construct significant legislative barriers to change. A commitment that comes to be regretted by future policymakers – applying the same set of principles as those applied originally (i.e. selecting from the same choice set) – is presumably more likely to be reversed than one that retains contingent support. In itself this may provide justification for incorporating institutional time consistency into the design process.

4 A model of capital taxation

Our focus for the rest of the paper will be on an environment that remains deterministic, but now features state variables. The aim is to generalise – and characterise more thoroughly – the idea of a time-consistent institution developed in the previous example. To do so when state variables are present is non-trivial. In particular, preferences across the set of constant promises are unlikely to have the time-consistency property. The best constant promise to implement in perpetuity when starting with a low capital stock is not generally the same as when starting with a high capital stock – and the capital stock will clearly evolve over time regardless of the chosen institution. Nonetheless, we show that choice sets exist in which time consistency can be achieved, and that by placing quite mild restrictions on the structure of these sets, much can be said about the character of time-consistent policy – particularly in steady state.

The particular example that we study is a variant of the optimal capital tax problem studied by Chamley (1986) and Judd (1985). It is a setting in which Ramsey policy delivers famously extreme outcomes, with initial capital taxes set to confiscatory levels – exploiting the inelasticity of the period-0 capital stock – but then with zero capital taxes in steady state. These outcomes are so asymmetric over time that it is not plausible to expect a policymaker ever to commit to a full Ramsey-optimal plan. As mentioned in the introduction, some authors have chosen to focus their policy advice on the zero capital tax steady-state result that follows from Ramsey policy. This is true, for instance, of Atkeson, Chari and Kehoe (1999). But we saw in the previous example that the steady-state outcome from a Ramsey plan may not have particularly desirable properties when it is considered in isolation. If time-non-contingent policy proposals are desired, it would be far preferable to obtain these through a deliberate, systematic normative procedure. This will, we hope, be a useful step towards increasing the value to policymakers of models of dynamic taxation.

4.1 Setup

The model is a deterministic representative-agent dynamic tax problem, in which the government each period must choose linear taxes on capital and labour income in order to finance a fixed expenditure requirement, which for simplicity we assume to be constant over time. The representative agent obtains positive utility from a single consumption good, denoted c_t for period t, and negative utility from supplying labour, l_t. The agent ranks dynamic consumption and labour sequences from period 0 onwards according to the function U_0, which takes the usual time-separable form:

$$U_0 := \sum_{t=0}^{\infty} \beta^t u\left(c_t, l_t\right) \tag{16}$$

with $\beta \in (0,1)$ and u satisfying the Inada conditions. The government is benevolent, and will treat U_0 as its social welfare function from the perspective of period 0.

In order to guarantee uniqueness of the time-consistent steady state in what follows, we must impose a further technical condition on preferences. The mathematical statement of this is not intuitive, and we relegate it to Section 7.1 in the appendix. In words, it says that if an agent's consumption and labour supply are increased linearly together, then the wage rate that would be needed to support the associated allocation must increase at an increasing rate. It is easy to show that the condition is satisfied by standard isoelastic preferences:

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

provided $\sigma \geq 1$ and $\nu \geq 1$ are jointly true. This is consistent with conventional assumptions on the Frisch and intertemporal substitution elasticities. Our arguments can be adjusted to accommodate preferences that are not restricted in this way, but this complicates the exposition with little gain.

The net production of goods takes place each period according to the fixed function $F\left(k, l\right)$, given capital inputs k and labour inputs l. F exhibits the usual properties.[18] Firms choose labour and (rented) capital inputs each period, given the gross wage rate w_t and capital rental rate r_t, both expressed in units of production. Hence they will set $w_t = F_{l,t}$ and $r_t = F_{k,t}$ in order to minimise their production costs. Capital is owned by consumers, and depreciates at fixed rate δ. Consumer income from capital, net of a depreciation allowance, is taxed at rate τ_t^k in period t, and labour income is taxed at rate τ_t^l. Hence the consumer's 'flow' budget constraint for period t is:

$$c_t + k_{t+1} \leq \left(1 - \tau_t^l\right) w_t l_t + \left[\left(1 - \tau_t^k\right)\left(r_t - \delta\right) + 1\right] k_t \tag{17}$$

Every period the government must be able to fund its fixed consumption requirement, $g > 0$. For the problem to deliver a non-trivial time-inconsistency problem, some restriction must be imposed that prevents this from being done by levying a sufficiently high, one-off capital tax in the initial time period, and zero taxes thereafter. In principle this is a feasible strategy, as it can allow the government to obtain a large positive net asset position for period 1 onwards, and fund all future expenditure from the interest on these assets. Because the initial capital stock is perfectly inelastic, this scheme is completely non-distortionary, and obtains a first best on the set of resource-feasible allocations. It is also fully time consistent, as continuing to fund expenditure from the

[18] That is, it is C^1, weakly concave, and exhibits constant returns to scale.

government asset stock is always an optimal continuation strategy once the initial capital levy has been applied. One approach to ruling this out is to impose an upper bound on the capital tax $\overline{\tau}^k$, or to assume that the period-0 capital tax is fixed *ex ante*.[19] We will take the alternative approach of assuming that the government must run a balanced budget each time period. This has the advantage of keeping the state vector unidimensional, which makes the dynamics that result from alternative institutional designs relatively easy to interpret. The government budget constraint for period t is then given by:

$$g \leq \tau_t^l w_t l_t + \tau_t^k \left(r_t - \delta\right) k_t \tag{18}$$

Given an initial capital stock k_0, a *competitive equilibrium with taxes* is a sequence of real allocations $\{c_t, l_t, k_{t+1}\}_{t=0}^{\infty}$, factor prices $\{w_t, r_t\}_{t=0}^{\infty}$ and taxes $\left\{\tau_t^l, \tau_t^k\right\}_{t=0}^{\infty}$ such that: (a) the real variables solve the consumer problem of maximising U_0 on the sequence of flow budget constraints, (17), and transversality; (b) firms are behaving optimally, with $w_t = F_{l,t}$ and $r_t = F_{k,t}$ for all t; (c) the government budget constraint (18) holds for all t; and (d) the following aggregate resource constraint is satisfied:

$$c_t + k_{t+1} + g \leq F\left(k_t, l_t\right) + \left(1 - \delta\right) k_t \tag{19}$$

The basic policy problem is to select a desired equilibrium from within the set of competitive equilibria with taxes. The following proposition simplifies this problem substantially. The proof uses arguments that are well-known in the dynamic Ramsey taxation literature (e.g. Chari and Kehoe, 1999), and is therefore omitted.

Proposition 1 *An interior allocation $\{c_t, l_t, k_{t+1}\}_{t=0}^{\infty}$ can be supported as a competitive equilibrium with taxes iff for all $t \geq 0$ it satisfies (a) the resource constraint (19) and (b) the following 'implementability' restriction:*

$$\beta^{-1} u_{c,t} k_{t+1} \leq u_{c,t+1}\left(c_{t+1} + k_{t+2}\right) + u_{l,t+1} l_{t+1} \tag{20}$$

Hence the implementability constraint is sufficient to ensure prices and taxes can be found that will support the given real allocation as a competitive equilibrium with taxes. The policy problem can then be expressed as direct choice over real allocations alone – the so-called 'primal approach' to optimal taxation. In words, the implementability condition states that the consumer must be on his or her flow budget constraint between period t and period $t+1$, given a set of prices (wages and effective real interest rates) that are consistent with marginal rates of substitution perceived in period t. The constraint applies only in period t: it is not a structural restriction in $t+1$. In general the government will have an incentive to violate it in period $t+1$, taxing away the consumer's capital income. Hence it is the implementability constraint that gives the problem a Kydland and Prescott character.

[19] Atkeson, Chari and Kehoe (1999) take the former approach, and Chari and Kehoe (1999) the latter.

4.2 A two-part problem

As with the example in Section 3, we can divide the normative problem of choosing a dynamic commitment strategy into an institutional design component and a 'day-to-day' policymaking component. To do so we again make heavy use of promise values. Here these correspond to promises that the total market value of the consumer's expenditure in a given period, less the value of their labour income – that is, the object on the right-hand side of (20) – will equal some minimum value. Again, we can divide the relevant constraint into promise-making and promise-keeping restrictions, requiring a dynamic sequence of promise values $\omega^{0,\infty}$ that satisfy, for all t:[20]

$$u_{c,t}k_{t+1} \leq \beta\omega_{t+1} \tag{21}$$

$$\omega_t \leq u_{c,t}\left(c_t + k_{t+1}\right) + u_{l,t}l_t \tag{22}$$

4.2.1 Day-to-day choice

Given $\omega^{0,\infty}$, the *day-to-day problem* is:

$$\max_{\{c_t,l_t,k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u\left(c_t, l_t\right)$$

subject to (19), (21) and (22), and a given initial capital stock k_0. So long as the commitment to $\omega^{0,\infty}$ is fixed, this problem is time consistent. Placing multiplier $\beta^t\mu_t$ on the resource constraint (19), $\beta^t\lambda_t^1$ on the promise-making constraint (21), and $\beta^t\lambda_t^2$ on the promise-keeping constraint (22), necessary conditions for an interior optimum are:

$$\begin{aligned}
0 &= u_{c,t} - \mu_t - \lambda_t^1 u_{cc,t}k_{t+1} + \lambda_t^2\left\{u_{c,t} + u_{cc,t}\left(c_t + k_{t+1}\right) + u_{cl,t}l_t\right\} && (23)\\
0 &= u_{l,t} + F_{l,t}\mu_t - \lambda_t^1 u_{cl,t}k_{t+1} + \lambda_t^2\left\{u_{l,t} + u_{cl,t}\left(c_t + k_{t+1}\right) + u_{ll,t}l_t\right\} && (24)\\
0 &= -\mu_t + \beta\mu_{t+1}\left[1 + F_{k,t+1} - \delta\right] - \lambda_t^1 u_{c,t} + \lambda_t^2 u_{c,t} && (25)
\end{aligned}$$

We denote by $V_0^d\left(\omega^{0,\infty}, k_0\right)$ the optimised value of the day-to-day problem from the perspective of period 0.[21]

4.2.2 Ramsey institutional design

The *institutional design problem* is then to choose a promise sequence $\omega^{0,\infty}$ according to some set of normative criteria. A Ramsey-optimal institution chooses $\omega^{0,\infty}$ to maximise $V_0^d\left(\omega^{0,\infty}, k_0\right)$ on the unrestricted domain $\mathbb{R}^\infty$, taking k_0 as given. Applying the evelope condition as before, necessary conditions are:

$$\lambda_t^1 = \lambda_{t+1}^2 \tag{26}$$

[20] We retain notational conventions from Section 3 regarding promise sequences.

[21] It is possible that the constraint set implied by (19), (21) and (22) is empty for the given $\omega^{0,\infty}$. In this case we set $V_0^d\left(\omega^{0,\infty}, k_0\right) = -\infty$.

for all $t \geq 0$ and:

$$\lambda_0^2 = 0 \tag{27}$$

Substituting (26) into (25) and imposing that all variables are constant, one recovers the famous Chamley-Judd result that steady-state capital taxes should be zero under Ramsey policy. In particular, the implication for steady state is:

$$1 = \beta \left[1 + F_{k,ss} - \delta\right] \tag{28}$$

which is only consistent with optimising behaviour on the part of consumers if $\tau_{ss}^k = 0$.

Condition (27), by contrast, confirms that in period 0 the policymaker is unencumbered by any past commitments to guarantee the consumer a particular level of within-period wealth, and will therefore be free to levy a very high capital tax. Figure 4 shows that this is indeed what is done. It charts the evolution of the capital tax rate τ_t^k under the Ramsey-optimal plan, given a standard parameterisation.[22] The units are simple proportions of net capital income, so the Figure shows a Ramsey-optimal capital tax in period 0 of the model that exceeds 300 per cent of net capital income.[23] The tax rate then decays to zero gradually over time.

Figure 5 illustrates the capital dynamics that accompany this tax policy. This highlights very starkly the time inconsistency problem: despite starting at its ultimate steady state value, the high initial tax rates cause capital to drop by around 5 per cent over the first ten years of the Ramsey plan, only recovering gradually to steady state over the course of the subsequent decades.

4.2.3 Ramsey policy: discussion

In generating Figure 4 we have assumed an initial capital stock k_0 that coincides with the long-run steady-state k_{ss}. This reinforces the time-contingent nature of Ramsey policy. In period 0 the economy is in almost exactly the same state as it is in period 60, and yet tax policy could scarcely be more different between the two cases.

Even if policymakers had complete confidence in the present model as a positive device, it does not seem realistic that they would ever be willing to commit to such a plan. The asymmetry in its prescriptions over time renders it an implausible choice. Precisely why it is implausible is hard to isolate. One

[22] Specifically, we assume a constant-elastic preference specification:

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

with $\sigma = 1$ and $\nu = 2$. A period is taken to be one year, and the discount factor β is, accordingly, set to 0.96. The production function F is a standard Cobb-Douglas with capital share of 0.33, and the depreciation rate δ is 0.05. Finally, g is set to 0.6, which implies that government consumption is equal to 34 per cent of net output in steady state.

[23] Clearly one may wish to place an upper bound on capital taxes at 100 per cent (or lower) for plausibility reasons, but there is no economic reason why obligations in excess of the legal 'tax base' could not be imposed.

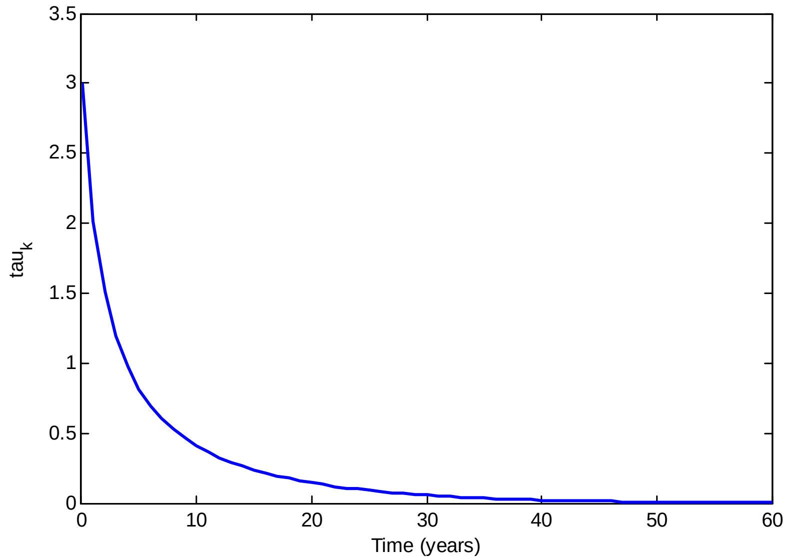

Figure 4: Ramsey policy: path of τ_t^k when $k_0 = k_{ss}$.

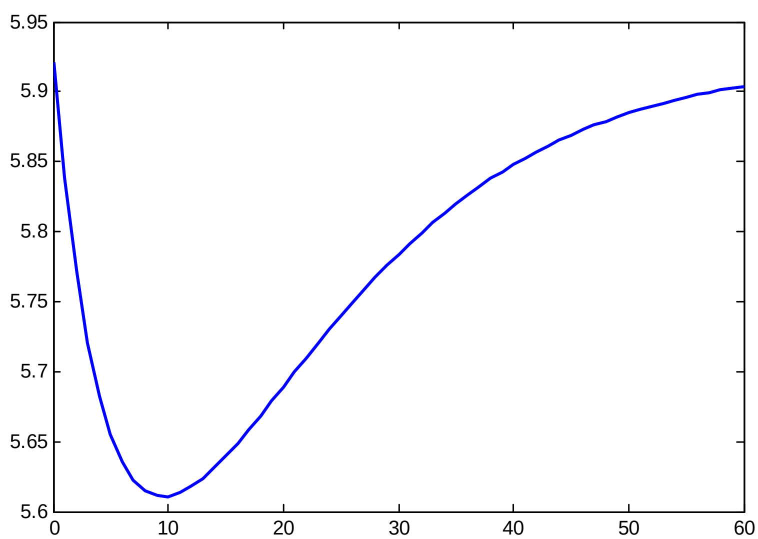

Figure 5: Ramsey policy: path of k_t when $k_0 = k_{ss}$

possibility is that time-contingent commitments are difficult to rationalise to the public. It may be easy to explain why capital taxes should change with the state of the economy, but less simple to explain why they should be set to punitive levels simply because the current year has been decreed 'period 0' in the model.

Alternatively, it may be that commitments which exploit the absence of any initial obligations – as implied by condition (27) – and have a time-contingent character as a consequence, are not considered credible. A popular maxim is that 'actions speak louder than words'. One way to establish credibility for a new institution in the first year of its existence is to demonstrate immediately a readiness to adhere to self-imposed constraints. But if the institution prescribes a period-0 policy such that ω_0 is not binding, then this cannot be achieved.

Finally, there are normative arguments that may count against time-contingent policies. The Ramsey-optimal outcome imposes the preference ranking of the period-0 policymaker, unconstrained, on all time periods. Even if it is known that subsequent policymakers will find continuation with the Ramsey institution an abhorrence, these later preferences are not considered. It is perfectly reasonable to have ethical doubts about this approach. Why should preferences from one time period alone be treated so differently from all others? Is it fair on future generations, or even democratic, to require that they should implement an allocation that they had no scope to choose?

Whatever the ultimate reason, it does seem that an economist who recommends the Ramsey-optimal institution is unlikely to have much practical influence. If commitment policy in Kydland and Prescott environments is to be founded on sound analysis, an alternative normative perspective on institutional design may be the best way to bridge the gap between theory and practice.

5 Time-consistent institutional design

In the simple two-period OLG model of Section 3 there was an obvious way to specify a restricted institutional design problem in which choice would be time-consistent. Since the environment was completely stationary, preference orderings over constant promise sequences were bound to remain invariant over time. But with capital accumulation, preferences over constant promise sequences may easily change. Nonetheless, we show in what follows that there will still exist alternative restricted sets of feasible institutions from which policymakers in all time periods can agree on an optimal choice. If choice is restricted – by convention or by constitutional mandate – to these sets, the resulting institution will have the appealing normative property that its selection is not regretted *ex post*, given the same set of options. Again, this is not time consistency in the 'discretionary' sense of allowing repeated, unrestricted choice over allocations. Promise-keeping constraints may still bind under the chosen institution. Time consistency instead applies to the selection process for the institutions that constrain allocations.

5.1 Preferences and time consistency

As in Section 3, our focus is on consistency between the preference orderings of policymakers at different points in time, ranking alternative dynamic promise sequences. The added complication now is that state variables will influence these preference orderings. Conditional on k_0, the period-0 value function $V_0^d\left(\omega^{0,\infty}, k_0\right)$ generates a binary preference relation over $\mathbb{R}^{\infty}$ that we denote $\succsim_{0,k_0}$. As before, this period-0 preference ordering can be used to rank promise sequences from generic period t onwards, $\omega^{t,\infty}$, based on the function $V_0^d\left(\left\{\widehat{\omega}^{0,t-1}, \omega^{t,\infty}\right\}, k_0\right)$ where for each $\omega^{t,\infty}$:

$$\widehat{\omega}^{0,t-1} \in \arg \max_{\{\omega^{0,t-1}:\{\omega^{0,t-1},\omega^{t,\infty}\}\in\Omega_0\}} V_0^d\left(\left\{\omega^{0,t-1}, \omega^{t,\infty}\right\}, k_0\right)$$

given a period-0 institutional constraint set Ω_0. This ranking is denoted $\succsim_{0,t,k_0}$, and is defined on the set Ω_t, constructed as before as the set of promise sequences from t onwards admitted under Ω_0:

$$\Omega_t = \left\{\omega^{t,\infty} \in \mathbb{R}^{\infty} : \left\{\omega^{0,t-1}, \omega^{t,\infty}\right\} \in \Omega_0 \text{ for some } \omega^{0,t-1} \in \mathbb{R}^t\right\}$$

The rankings $\succsim_{t,k_t}$ and $\succsim_{t,s,k_t}$ $(s > t)$ are defined analogously over $\mathbb{R}^{\infty}$ and Ω_s respectively.

We can then define time consistency in the institutional design problem as follows:

Definition 1 *For the constraint set* Ω_0 *and given* k_0*, let* $\widehat{\omega}^{0,\infty} \in \arg\max_{\{\omega^{0,\infty}\in\Omega_0\}} V_0^d\left(\omega^{0,\infty}, k_0\right)$*, and suppose that the series* $\{k_t^*\}_{t=0}^{\infty}$ *is an optimal capital choice for the day-to-day policymaker given* k_0 *and* $\widehat{\omega}^{0,\infty}$*.*[24] *We say that institutional choice is* ***time consistent in*** Ω_0 *if for all* $t > s \geq 0$*, the rankings* $\succsim_{t,k_t^*}$ *and* $\succsim_{s,t,k_s^*}$ *over* Ω_t *coincide.*

Note that this is an 'on-equilibrium' definition of time consistency. That is, it requires consistency between the preferences of initial policymakers and their successors only along the capital accumulation path that is induced when the optimal institution in Ω_0 is selected. The preference mapping $\succsim_{t,k_t}$ generally differs in k_t, and so it is not possible to guarantee that $\succsim_{t,k_t}$ will coincide with $\succsim_{s,t,k_s}$ for arbitrary k_t. Nonetheless, it captures our idea of time-consistent institutional design as the absence of regret. Assuming that day-to-day choice evolves in a manner consistent with the optimal institution in Ω_0, the decision to select that institution – given the same choice set for t onwards – will not be challenged.

The definition does not rule out trivial statements of time-consistent choice: for instance, institutional choice is time-consistent in any Ω_0 set that contains just a single feasible institution, and this institution could always be the period-0 Ramsey plan. In this sense one could conceivably justify the Ramsey policy as a time-consistent one in the set that contains it alone – but this is not

[24] We write $k_0 \equiv k_0^*$ for simplicity.

a particularly useful claim. We are instead interested in the consequences of reestricting Ω_0 in as parsimonious a manner as possible. In particular, what are the minimal restrictions that need to be placed on Ω_0 to obtain time consistency in the institutional design problem? These restrictions prevent institutional designers from taking any decisions that subsequent generations come to regret. They can perhaps best be thought of as constitutional restrictions on the manner in which institutions can be designed, whilst the act of choosing institutions within Ω_0 can be thought of as the direct legislation process that fixes long-term commitments.

It turns out that we are able to go quite far in describing what a 'minimally restricted' set Ω_0 looks like. In doing so the idea of a 'degree of freedom' for policy choice is very useful. We mean by this the following:

Definition 2 *Consider a set of possible institutions $\Omega_0 \subseteq \mathbb{R}^\infty$. We say that Ω_0 exhibits n degrees of freedom at $\widehat{\omega}^{0,\infty} \in \Omega_0$ if there exist exactly n linearly independent subspaces of $\mathbb{R}^\infty$ along which the promise series can be varied differentially at $\widehat{\omega}^{0,\infty}$ whilst remaining in Ω_0.*

Thus the number of degrees of freedom at a given point in Ω_0 is the number of distinct, joint changes to promise values that the policymaker can make when choice is restricted to Ω_0. The unrestricted set $\mathbb{R}^\infty$ has an infinite number of degrees of freedom, as any ω_t can be chosen for period t without affecting what can be chosen in any other period. The set of constant promises has just one degree of freedom: it is impossible to vary ω_t except by varying all other promises by an equal amount. The singleton set that contains just the Ramsey plan has no degrees of freedom.

The following proposition then allows us to place substantial structure on choice sets that satisfy time consistency:

Proposition 2 *Suppose institutional choice is time-consistent in Ω_0, and let $\widehat{\omega}^{0,\infty}$ maximise $V_0^d\left(\omega^{0,\infty}, k_0\right)$ on Ω_0 given k_0, such that $V_0^d\left(\widehat{\omega}^{0,\infty}, k_0\right) > -\infty$. Then Ω_0 has at most one degree of freedom at $\widehat{\omega}^{0,\infty}$.*

Proof. See appendix. ■

Proposition 2 suggests that a requirement for time consistency in Ω_0 will simplify substantially the institutional design process: choice need be made only along one degree of freedom, at least in the neighbourhood of the optimum.[25] This means that choice within a time-consistent Ω_0 is effectively equivalent to selecting the value of at most one parameter, where promises for all time periods are continuously differentiable in this parameter. To understand this point better, notice that the proof of Proposition 2 could be applied without changes to the stationary problem of section 3. In that setting it implies that the

[25] More generally, it is a straightforward generalisation of the proof of Proposition 2 that a model whose structure implies an n-dimensional promise vector will have n degrees of freedom for time-consistent institutional choice. We are grateful to Ramon Marimon for highlighting this.

set of constant institutions admits as many degrees of freedom as it is possible to have in a choice set that is time-consistent. There is unidimensional parametric choice, where the unique parameter is 'the' promise value that will obtain in all periods.

5.2 Steady-state capital taxes

Importantly, it does not follow from Proposition 2 that there is just one possible parametric class of institutions in which choice is time-consistent: it is just that if there is a time-consistent choice set, it is unidimensional. To achieve a unique Ω_0 we will need to impose some restrictions the structure of admissible choice sets. It turns out, however, that a relatively weak set of restrictions is necessary in order to say something quite substantive about the outcomes of time-consistent choice *in steady state*. In particular, we show that three quite parsimonious assumptions are enough to tie down a unique steady-state outcome. Importantly, and in direct contrast with the steady state of a Ramsey model, capital taxes will be positive in this steady state.

5.2.1 Three restrictions on Ω_0

We now explain in turn the three assumptions on Ω_0 that we impose. First, we know from Proposition 2 that a minimal necessary restriction is to limit choice to a single dimension in the neighbourhood of the optimum, and without any loss we can impose that there is a single degree of freedom everywhere in Ω_0. This motivates a focus exclusively on classes of Ω_0 that have a unidemensional parametric form, and our first assumption formalises this. It restricts attention to Ω_0 for which there is continuous differentiability of the promises in some parameter θ. Hence there is one degree of freedom for choice everywhere in Ω_0. Formally:

A 1 *(Parametric Ω_0) There exists a convex set $\Theta \subseteq \mathbb{R}$ and array of C^1 functions $\{\omega_t : \Theta \to \mathbb{R}\}_{t=0}^{\infty}$ such that:*

$$\Omega_0 = \left\{\omega^{0,\infty} \in \mathbb{R}^\infty : \omega^{0,\infty} = \{\omega_t(\theta)\}_{t=0}^{\infty} \text{ for some } \theta \in \Theta\right\}$$

Subject to this unidimensional form – which is necessary for time consistency – we wish to restrict Ω_0 in the minimal way possible. Our second assumption is a way to ensure this. It requires that any value for ω_t in $\mathbb{R}$ could be obtained by appropriate choice in Ω_0. This assumption prevents arbitrary boundaries being imposed on Ω_0 in a way that would ensure some preferred choice – ruling out, for instance, restricting Ω_0 to a singleton that contains just the period-0 Ramsey plan. Formally:

A 2 *(Full domain) For any $t \geq 0$ and $\omega_t' \in \mathbb{R}$ there exists some $\widehat{\omega}^{0,\infty} \in \Omega_0$ such that $\widehat{\omega}_t = \omega_t'$.*

A final assumption provides a bit more structure on the possible variation of choice within Ω_0. It ensures that any movement within the set of possible institutions cannot induce unboundedly large changes in one period's promises relative to another's. It strengthens A2, which would be consistent, for instance, with ω_t varying by an amount β^t for every unit by which ω_0 is varied. Clearly as $t \to \infty$ this eliminates the scope for variation in ω_t except at a punitive cost in terms of changes to the earlier promises. To prevent this, and ensure meaningful scope for choice at every horizon, we impose:

A 3 *(Bounded relative variation) For any scalar $\varepsilon > 0$ there exists a scalar $\delta > 0$ such that for all $\widehat{\omega}^{0,\infty}, \widetilde{\omega}^{0,\infty} \in \Omega_0$, $|\widehat{\omega}_t - \widetilde{\omega}_t| < \varepsilon$ implies $|\widehat{\omega}_s - \widetilde{\omega}_s| < \delta$ for all $s, t \geq 0$.*

5.2.2 Deterministic steady state

Assumptions A1 to A3 are certainly non-trivial, but we view them as quite mild in character. They simply ensure that there is meaningful choice across institutions in a set that has the maximum number of degrees of freedom possible for a time-consistent choice set. Despite this, together they are sufficient to ensure a very strong result on steady-state outcomes.

Proposition 3 *Suppose that choice within a set Ω_0 is time-consistent, and that Ω_0 satisfies assumptions A1 to A3. Then there is a unique steady state to which real allocations can converge, defined independently of the precise specification of Ω_0. This steady state features positive capital taxes.*

Proof. See appendix. ■

Thus time consistency *per se* is enough to determine steady-state outcomes, without any additional restrictions on Ω_0 needed. Under assumptions A1 to A3, the steady state of any time-consistent institutional choice will be characterised by the solution to a set of non-linear simultaneous equations, with as many equations as there are variables to solve for. The maintained assumption on preferences given by Condition 1 in the appendix ensures that these equations have a unique solution.

The table below compares the values of capital and labour income tax rates in steady state under Ramsey and time-consistent institutions, using the same parameterisation as Figure 4.

	Ramsey-optimal	Time-consistent
τ^k_{ss}	0	0.28
τ^l_{ss}	0.54	0.50

The absolute magnitude of the increase in capital taxes between the Ramsey and time-consistent steady states – 28 percentage points – is clearly much greater than the associated reduction in labour income taxes. But one should be cautious about dismissing the merits of the time-consistent allocation on this basis alone. Assuming that welfare losses are roughly proportional to the sum of the squares of the taxes, there is surprisingly little to separate the two. A direct welfare comparison between the two steady states will find in favour of the Ramsey path, but this is a misleading comparison as the capital stock is lower in the time-consistent case. This implies lower welfare, but only after different transition paths have been followed. Steady-state welfare is not the main policy criterion: if it were then the Ramsey policy could be improved upon still further by letting the policymaker's discount factor approach unity.

Why does time-consistent policy involve positive steady-state taxes? The reason is that time-consistent institutional choice must be made jointly for all time periods, along one parametric dimension. This means that any benefits from relaxing the period-0 promise-keeping constraint must be accompanied by costs from tightening the period-0 promise-making constraint – along with additional changes to these constraints at all future horizons. This contrasts with the Ramsey institutional design problem, in which benefits from relaxing the period-1 promise-keeping constraint are equated with costs from tightening the period-0 promise making constraint. The difference in timing implies a difference in relative discounting, with the effect that instead of the steady-state multipliers satisfying $\lambda_{ss}^1 = \lambda_{ss}^2$, as implied by (26) in the Ramsey case, we have $\beta\lambda_{ss}^1 = \lambda_{ss}^2$.

Very loosely, the effect is a policy that strikes a balance between the short-term incentive for high capital taxes and the long-run incentive to eliminate distortions. Recall that the main reason for zero steady-state capital taxes being optimal in the Ramsey case is that a positive rate would imply an effective tax wedge on the period-t good that is growing at a compound rate in t. This is generally inefficient relative to a constant wedge. But the losses that result from departing from zero steady-state taxes do not accrue immediately, and the lower is β the more they will be discounted when institutions are initially chosen. By contrast, there are gains to allowing high capital taxes in the very first time period, since capital is inelastic initially. The joint restrictions on choice that are required by time consistency mean that high taxes can only be implemented initially if they are also implemented in steady state. Hence the long-term losses from positive steady-state taxes are traded off against short-term gains from taxing an inelastic factor. The result is a positive – but modest – steady-state capital tax. Only as $\beta \rightarrow 1$ is the optimality of zero steady-state taxes restored.

For this reason it can certainly be argued that positive steady-state capital taxes are the product of a 'short-termist' perspective on policy. Any gains relative to a zero capital tax policy will be exhausted during the early periods of the model. But if a more long-term perspective on optimality is desired then, again, the obvious way to achieve it is to raise the policymaker's discount factor above the household's β. So long as the future is discounted it should not be surprising that high capital taxes are favoured, given the trade-off we have constructed.

6 Transition dynamics

Proposition 3 provides extremely useful insight into the long-run character of time-consistent institutions, but it says nothing about transition paths. For this we need to place more structure on the sets from which time-consistent choice is made. In general there are many Ω_0 sets satisfying assumptions A1 to A3 in which choice is time-consistent. Each of these will be associated with a distinct optimal promise sequence, and hence a distinct prescription for optimal taxes. Some additional selection is thus needed. In this section we propose three further restrictions on the character of optimal policy, and show that these are sufficient to deliver a unique choice. Whilst the resulting allocations have appealing properties, our approach is just one possible way to resolve the multiplicity. We remain open to the possibility that alternatives may have equal or greater appeal, and clarifying this issue will be an interesting avenue for future work.

6.1 Obtaining uniqueness

The problem we have is how to select a unique time-consistent choice set Ω_0, given the initial capital stock k_0. The set Ω_0 must vary in k_0 if it is to be time-consistent for all k_0. If it did not then the only clear candidate satisfying A1 is the set of constant promises, but it is easy to show that choice within this set is not time-consistent, so long as k_0 is away from the steady state characterised in the previous section.[26] Thus we are searching for a unique set-valued mapping $\Omega_0(k_0)$, with $\Omega_0 : \mathbb{R} \to \mathbb{R}^\infty$, where the output of this mapping must satisfy A1 to A3.

Associated with each $\Omega_0(k_0)$ is an optimal institutional choice, $\widehat{\omega}^{0,\infty}(k_0)$, which satisfies:[27]

$$\widehat{\omega}^{0,\infty}(k_0) \in \arg\max_{\omega^{0,\infty} \in \Omega_0(k_0)} V_0^d\left(\omega^{0,\infty}, k_0\right) \tag{29}$$

In response to the chosen institution $\widehat{\omega}^{0,\infty}(k_0)$, capital evolves optimally according to the series $\{k_t^*(k_0)\}_{t=1}^\infty$. By definition, if choice in $\Omega_0(k_0)$ is time-consistent then the continuation of $\widehat{\omega}^{0,\infty}(k_0)$ from t onwards, $\widehat{\omega}^{t,\infty}(k_0)$, will then be an optimal choice under preferences $\succsim_{t,k_t^*(k_0)}$ in the corresponding continuation set $\Omega_t(k_0)$, defined as before for all k_0:

$$\Omega_t(k_0) = \left\{\omega^{t,\infty} \in \mathbb{R}^\infty : \left\{\omega^{0,t-1}, \omega^{t,\infty}\right\} \in \Omega_0(k_0) \text{ for some } \omega^{0,t-1} \in \mathbb{R}^t\right\}$$

The additional restrictions that we make will relate to the structure of $\Omega_0(k_0)$ and $\Omega_t(k_0)$ as k_0 varies.

[26] In principle time-varying unidimensional choice sets – such that Ω_t differs from Ω_0, but both are invariant to k_0 – would also be consistent with A1. Again, it is easy to show that these cannot deliver time-consistent choice for all k_0.

[27] It is possible that the argmax set may contain multiple elements. In this case it is unlikely that more than one of these implies time-consistent choice, and $\widehat{\omega}^{0,\infty}(k_0)$ is then the unique element that does so. If there is more than one element in the argmax set that implies time-consistent choice, selection among them can be by any arbitrary procedure.

6.1.1 Two further restrictions on Ω_0

Our first step towards obtaining uniqueness is the following assumption:

A 4 *(Time invariance) For all $k_0 \in \mathbb{R}_+$ and all $t > 0$:*

$$\Omega_t(k_0) = \Omega_0(k_t^*(k_0)) \tag{30}$$

That is, the continuation set of promise series available in Ω_t when the initial capital stock was k_0 must coincide with the initial set of promise series available when the initial capital stock is $k_t^*(k_0)$. Put differently, there can be no time contingency in the choice set: if the capital stock in period t is k_t then there is a unique possible structure for Ω_t, whether $t = 0$ and Ω_t is the initial choice set, or $t > 0$ and Ω_t is a continuation of Ω_0, with $k_t = k_t^*(k_0)$ for an initial k_0. It is, above all, the time-contingent character of Ramsey policy that makes it an implausible candidate for practical policy design. Since our aim is to design a choice framework that will deliver more realistic rules than the Ramsey framework, we lose little from ruling out such contingency.

The second assumption is the following:

A 5 *(No envy) Fix $k_0, k_0' \in \mathbb{R}_+$, and suppose the series $\omega^{0,\infty}(k_0) \in \Omega_0(k_0)$ and $\omega^{0,\infty}(k_0').\in \Omega_0(k_0')$ converge to the same $\omega \in \mathbb{R}$. Then $\omega^{0,\infty}(k_0) \succsim_{0,k_0} \omega^{0,\infty}(k_0')$.*

An $\Omega_0(k_0)$ choice set that satisfies assumptions A1 to A4 is one in which there is one dimension of choice, and in which (by A2) every $\omega \in \mathbb{R}$ is a possible selection for the long-run promise value. This means that choice over $\Omega_0(k_0)$ is equivalent to choice of the long-run promise. To the extent that the sets $\Omega_0(k_0)$ and $\Omega_0(k_0')$ differ, it is in the structure of their promise sequences prior to convergence. Assumption A5 considers promise sequences that induce convergence to the same long-run promise value – and in this sense are 'equivalent' choices – for distinct k_0. The assumption requires that among any set of promise sequences that are equivalent in this sense, each policymaker should prefer the one that is actually available to them, relative to those that would be available for alternative k_0 values.

This is the stronger of the two assumptions made in this Section. It captures the idea that the set of options available to any one policymaker should not be envied by any other. Why impose such a condition? Variations in the initial capital stock will be associated, in general, with variations in the value of an optimal institution. This happens for two reasons. The first is the usual fact that more capital expands the feasible set of real allocations. The second is that variations in the initial capital stock induce changes in $\Omega_0(k_0)$, and this in turn affects the value of the policies available. This second factor may allow for some quite perverse outcomes. It is quite possible, for instance, to construct time-consistent choice sets satisfying assumptions A1 to A4, in which lower values for the capital stock are associated with a higher value for the optimal institution. This happens if lower k_0 values deliver an $\Omega_0(k_0)$ set in which the optimal initial promise allows for very high initial capital taxes, bringing outcomes close to the Ramsey plan, whilst higher values for k_0 are associated with low initial taxes.

There is something unsatisfactory about such cases. The variation in value induced by differences in the choice sets seems very arbitrary: some starting values for k_0 are simply given 'favourable' treatment relative to others. At the same time, ruling out all variation in the promise values *a priori* would be too strong. It is quite possible that when k_0 is low, a different promise sequence is best suited to delivering a desired steady state, relative to the case in which k_0 is high. A5 allows for differences between promise sequences that deliver the same steady state as k_0 varies, but only where these differences are mutually beneficial – that is, where an institutional designer in initial state k_0 will not envy the promises available in state k_0', and vice-versa.

6.1.2 Uniqueness

Given these additional restrictions, we can state the following. The proof is in the appendix.

Proposition 4 *For all $k_0 \in \mathbb{R}_+$, suppose that $\Omega_0(k_0)$ satisfies assumptions A1 to A5, that optimal choice in $\Omega_0(k_0)$ is time-consistent, and that real allocations converge to a steady state given this optimal choice. Then the optimal choice in $\Omega_0(k_0)$ sets $\widehat{\omega}_t(k_0) = \omega_{ss}^*$ for all t, where ω_{ss}^* is the unique steady-state promise consistent with Proposition 3.*

This result hinges on the 'no envy' assumption, A5. It means that any increase in k_0 at the margin can increase the value of optimal policy only through the 'direct effect' of a higher capital stock – there can be no 'indirect effect' due to a more (or less) favourable set of promise values. If there were, changing the promise values $\widehat{\omega}^{0,\infty}(k_0)$ by varying k_0 at the margin would raise the value of policy. This is inconsistent with the assumption that $\widehat{\omega}^{0,\infty}(k_0)$ is optimal in the set of possible $\widehat{\omega}^{0,\infty}(k_0')$ sequences when initial capital is k_0. But the only type of sequence for which the marginal benefit to raising the capital stock is always equal to the direct effect, irrespective of k_0, is a constant sequence – and the only constant promise consistent with time-consistent optimality and convergence is ω_{ss}^*.

6.1.3 Are constant promises too restrictive?

On the surface, imposing the same value for ω_t in every time period may appear unduly restrictive. It suggests that the market value of the consumer's net purchases should match a given, constant value each period, even whilst capital accumulation may be expanding the feasible set of allocations. But it is important not to neglect the units in which value is assessed here. Consider rewriting the promise-keeping constraint as follows:

$$\frac{\omega_t}{u_{c,t}} \leq \left[(c_t + k_{t+1}) - w_t\left(1 - \tau_t^l\right) l_t\right] \tag{31}$$

where we have substituted $w_t\left(1-\tau_t^l\right)$ in place of the marginal rate of substitution between consumption and labour supply. The object on the right-hand-side is value of the consumer's purchases net of labour income in period t – that is, purchases from wealth taken into period t – assessed in units of the real consumption good. If our restriction were that this object must be held constant, it would be an extremely strong one. In effect, the consumer's real disposable non-labour wealth would be being held constant, even as the aggregate capital stock of the economy might be growing. But a constant value for ω_t instead fixes the minimum value of these net purchases relative to the inverse of the marginal utility of consumption, $\frac{1}{u_{c,t}}$. This object, in turn, can be read as the marginal resource cost of an extra unit of utility in period t, likewise expressed in units of the consumption good. Hence the restriction fixes the real value of disposable non-labour wealth in period t only relative to real marginal cost of utility. In periods when real wealth is low, consumption will likewise be low, and so too will be the marginal cost of utility provision. Thus the constraint is sensitive to the dynamics of capital accumulation: it will be relatively easy to satisfy when the aggregate capital stock is low.

Another way to make the same point is to consider the promise-keeping constraint in the more general case that the optimal choice for ω_t varies in the contemporaneous value of, say, k_t, and so we can write: $\omega_t=\omega\left(k_t\right)$. In this case the promise-keeping constraint can be written:

$$\omega\left(k_t\right)\leq u_{c,t}\left(c_t+k_{t+1}\right)+u_{l,t}l_t \tag{32}$$

But evidently we could rewrite the constraint:

$$\eta_t\leq\frac{u_{c,t}\left(c_t+k_{t+1}\right)+u_{l,t}l_t}{\omega\left(k_t\right)} \tag{33}$$

The invariant choice of $\eta_t=1$ will clearly now be optimal for all t – the point being that re-basing the numeraire can always imply a constant 'promise' is optimal in some version of the problem. It just happens that under the 'no envy' assumption the version for which this is true is the one we study.

6.2 The structure of Ω_0

When these assumptions are made on $\Omega_0\left(k_0\right)$, the choice set can be studied numerically. Figures 6 and 7 provide two graphical representations of it, for cases in which capital starts 10 per cent below and above steady state respectively.[28] The lines in the figure correspond to alternative dynamic sequences for

[28] The figure is generated as follows. First we solve for the optimal dynamic path in Ω_0, given the (known) constant promise $\overline{\omega}$. Then we determine the loss to the present value of welfare that would follow if an alternative constant promise were chosen, starting with a capital stock equal to the time-consistent steady-state level. Finally we work backwards through time, finding a promise in each period that will deliver the same welfare loss in present value terms, given a starting capital stock consistent with the accumulation dynamic witnessed on the optimal path.

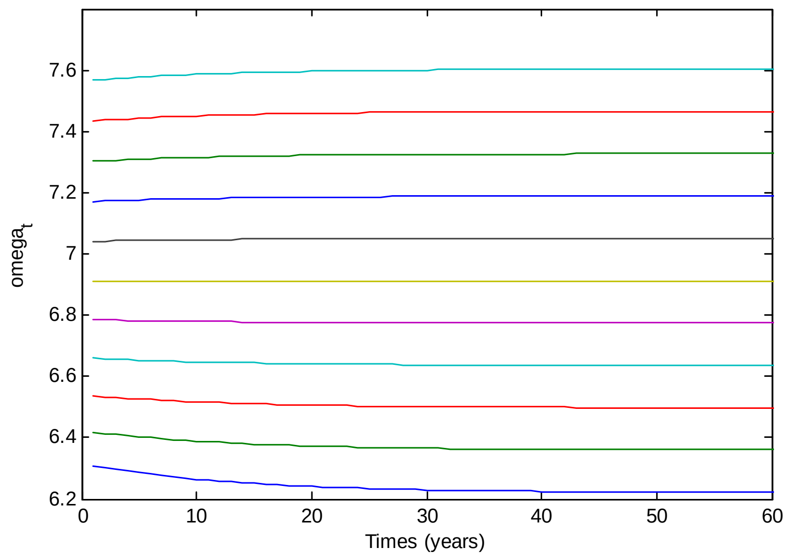

Figure 6: Structure of $\Omega_0(k_0)$: k_0 10 per cent below steady state

the promises, the institutional design choice being to select from among them.[29] In each case the optimal choice is the central, time-invariant line. Time consistency implies that this choice remains optimal as time progresses, given the associated capital accumulation dynamic. We compute suboptimal paths under the assumption that the loss to switching from the optimal path to any one of the alternatives graphed is invariant along the optimal path, where this loss is assessed at the point in time when the switch takes place.

Figures 6 and 7 illustrate the role of the capital stock on the shape of $\Omega_0(k_0)$. When capital starts below steady state, as in Figure 6, relatively small changes in initial promise values are able to induce relatively large changes in utility. The result is that the choice set bows inwards, as a given reduction in the value of the promise sequence is initially obtainable through smaller departures in the promises from the optimal choice. By contrast, when the capital stock starts high, as in Figure 7, greater initial departures in the promises are needed to obtain a given reduction in value: the choice set fans out. Once the capital stock is in steady state, time consistency is ensured by a time-invariant choice set: that is, the lines become parallel to one another.

The time-dependence of the sub-optimal choices in Figures 6 and 7 reflects the fact that optimal institutional choice in $\Omega_0(k_0)$ is time-consistent only along the equilibrium path. Suppose that in period s a sub-optimal institution in $\Omega_0(k_0)$ were to be selected, and suppose that optimal choice for the 'inner'

[29] The lines graphed are a selection from the continuum of elements in Ω_0.

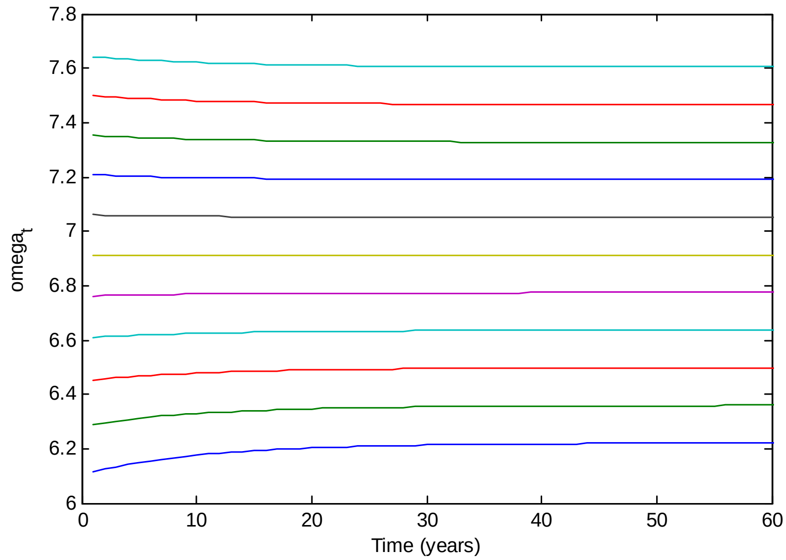

Figure 7: Structure of $\Omega_0(k_0)$: k_0 10 per cent above steady state

problem up to period $t > s$ then delivered a capital stock of k_t, given this instituion. Then the continuation choice set $\Omega_t(k_0)$ will not coincide with $\Omega_0(k_t)$ – that is, the time-consistent choice set associated with this unanticipated value of the capital stock. Optimal choice in the continuation set will thus no longer be time-consistent. The off-equilibrium paths in Figures 6 and 7 are tailored to the capital accumulation dynamic associated with optimal institutional choice, and any alternative accumulation will undermine the time-consistency of institutional preferences.

Our assumption throughout this paper is that a commitment technology is available, and for this reason the failure to guarantee time-consistent institutional design off-equilibrium is a second-order concern. Our aim is not to solve for the subgame-perfect equilibrium of a dynamic game: it is to devise a means for choosing long-run policy plans that will not later come to be regretted, given the same set of options.

6.3 Simulated dynamics

Consistent with the preceeding discussion, Figures 8 and 9 illustrate the dynamics of capital and labour tax rates when promise values are held fixed at their time-consistent steady-state level, and k_0 is set to 90 per cent, 100 per cent and 110 per cent of the time-consistent steady-state capital stock.

A low initial capital stock induces lower capital taxes, which incentivises

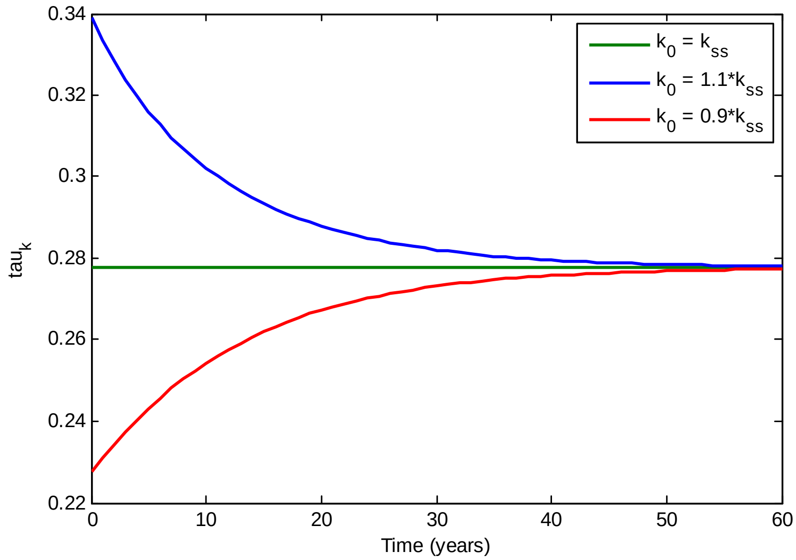

Figure 8: Capital taxes: time-consistent institutional policy

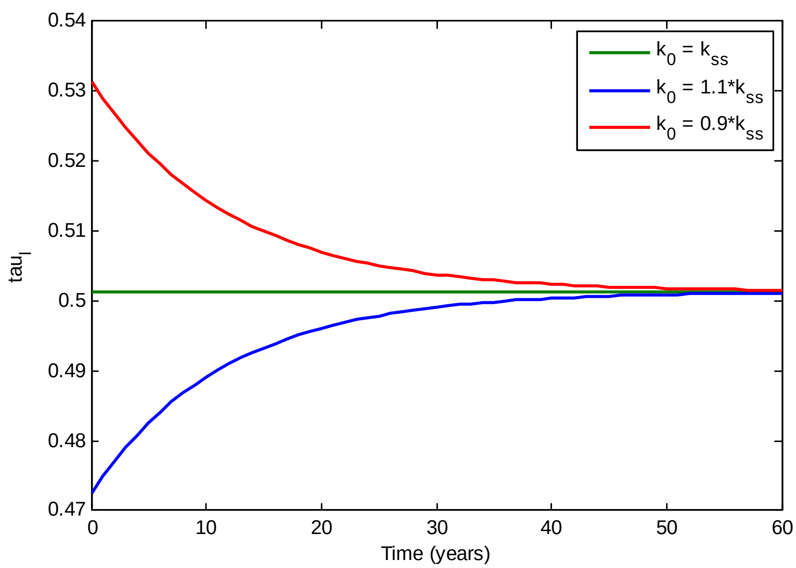

Figure 9: Labour taxes: time-consistent institutional policy

faster accumulation. Government expenditure requirements are fixed, and these lower capital taxes must be substituted with higher labour taxes initially. When capital starts above steady state the dynamics are roughly symmetric, with high initial capital taxes and low labour taxes. These qualitative patterns are not robust for all calibrations. Importantly, if the elasticity of intertemporal substitution falls below 0.5, a low initial capital stock comes to be associated with higher capital taxes and lower labour taxes, relative to steady state. A low value for the EIS corresponds to a situation in which the 'price effect' of lower consumption substantially dominates the 'quantity effect', which ultimately implies that the promise-keeping constraint can be relaxed by constraining capital accumulation – rather than enhancing it.

More significantly, it is clear that time-consistent institutional design does not share the time-contingent character of Ramsey policy – as distinct from state contingency. If the the economy starts in steady state, it will remain there in perpetuity. Recall, by contrast, that a Ramsey institutional designer who inherits a capital stock at the Ramsey steady-state level will raise capital taxes to punitive levels, and induce a fall in the capital stock before returning it gradually back to steady state.

6.3.1 Welfare comparisons

Comparing welfare between time-consistent institutions and alternatives – including Ramsey policy – is complicated by the fact that different capital dynamics are induced in each case. From any initial capital stock, the Ramsey-optimal policy is clearly welfare-superior, but the magnitude of the gains from it are surprisingly modest. If capital starts at the steady-state value associated with time-consistent policy, the welfare gains from a Ramsey institution relative to a time-consistent institution under our benchmark calibration are equivalent to a permanent consumption gain of 1.6 percentage points (from the time-consistent steady-state level). The variability of taxes is an order of magnitude lower under time-consistent policy, but welfare is certainly not.

An alternative comparison can be made by considering the consequences of switching to a time-consistent policy when starting in the Ramsey steady state – that is, assuming that many years have elapsed since Ramsey policy was first implemented, and considering the benefits to changing to an alternative constant set of promises. For our benchmark calibration the gains are positive, and equivalent to a permanent consumption gain of 0.4 percentage points relative to the Ramsey steady state. As emphasised in Section 5.2, this is not to say that the steady-state allocation associated with time-consistent policy is preferable to the steady-state outcome of Ramsey policy: the capital stock is ultimately lower in the former case, due to positive capital taxes. But if a society were to find itself in Ramsey steady state and consider the benefits of switching to a time-consistent institution, assessing welfare along the transition path, it would wish to do so.

The reverse does not apply: starting from steady state under the time-consistent institution and switching to the constant sequence of promises associated with Ramsey steady state will always be welfare worsening. Indeed, a direct corollary to Proposition 3 is that the steady state associated with optimal time-consistent policy is the only steady state for which there are no benefits from switching to an alternative constant promise sequence.[30]

The table below summarises these welfare comparisons. The first row lists the equivalent permanent percentage consumption gain associated with switching from the Ramsey steady state to, in turn, a full Ramsey-optimal path, and a time-consistent institution. The second row lists the equivalent consumption gain associated with switching from the time-consistent steady-state path to the Ramsey plan and the Ramsey steady-state institution.[31] In all cases these welfare gains are assessed at the point in time that the switch is made.

	New institution:		
k_0:	Ramsey	Ramsey ss	Time-consistent
Ramsey ss	1.91%	0	0.41%
Time-consistent ss	1.57%	−0.37%	0

7 Conclusion

When expectations of future policy affect what can be done today, time inconsistency exists in two distinct forms. The first is unavoidable: without some sort of institutional 'commitment device', policy variables will be determined as the outcome of a dynamic leadership game between policymakers at different horizons, each choosing only from the period-specific choice set available to them. It is well known that the equilibrium of this game need not be Pareto efficient: policymakers at all horizons would be better off under alternative feasible schemes. This is the sense in which 'time inconsistency' is widely understood as a concept, and large literatures exist contrasting outcomes in different problems when policy is set by 'commitment' and 'discretion'.

But a second form of time inconsistency also applies to the commitment problem itself, and this has been the main focus of the current paper. It is not just that all policymakers can be made better off by implementing 'rules rather than discretion'. The problem of designing appropriate rules – 'institutional design', as we have labelled it – is itself time-inconsistent, at least when there are no restrictions placed on the set of institutions that can be considered. The best promise for period t when choosing a Ramsey plan in period 0 differs from the best promise for period t when choosing a Ramsey plan in some other period $s \leq t$. Thus Ramsey plans would not survive an unexpected re-evaluation.

[30] It is possible that an alternative constant promise sequence could improve on the time-consistent institution when capital is away from steady state, but the choice of this alternative sequence would itself not be time consistent along the induced accumulation path. Recall that when k_0 departs from steady state, almost all of the promise sequences in Ω_0 are time-varying.

[31] The latter is interpreted as a switch to constant promises, equal to the promise value associated with Ramsey steady state.

Our purpose in this paper has been to ask what sorts of outcomes would result if the problem of designing institutions were restricted in a way that ensured it was time consistent. That is, suppose that dynamic sequences of policy commitments could only be chosen from a restricted set, deliberately constructed to ensure that the best choice from within this set in period 0 remains the best choice in a subsequent period s. What does optimal time-consistent choice of this kind look like? The idea is not that institutional design should be assumed actually to take place period-by-period. If this were true then we would return to the inefficient Stackelberg equilibrium of traditional 'discretion'. We are just interested in making sure that if an institution is re-compared with the same set of alternatives from which it was originally chosen, it should remain the best choice – always under the assumption that if it is chosen, it will endure permanently.

Applying this criterion to the institutional design problem can go a suprisingly long way in determining policy choice. In particular, we have shown in the context of an optimal capital tax problem that if institutional design is time-consistent then the choice set from which institutions are selected can have at most one degree of freedom in the neighbourhood of an optimum. This means that time-consistent choice must be equivalent to selecting a single 'policy parameter' that jointly determines commitments in all time periods. If the economic environment under study were stationary – with no endogenous state variables – then this choice parameter could be thought of as a constant policy rule, to be implemented in all time periods. When the state of the economy may evolve over time, the parameter is a single 'lever' for varying rules in all periods jointly, but the precise rule that is associated with a given period can also depend on the state of the economy.

More significantly, we have demonstrated that if institutions are required to be chosen from choice sets that admit time consistency, three apparently mild assumptions on the structure of these sets will be enough to induce a unique steady-state outcome.[32] This steady state differs in a systematic way from the steady state of Ramsey policy, most noticeably because it involves a positive, though relatively small, capital tax. The reason why a positive tax is optimal, despite implying that there will be long-run distortions to productive decisions in the economy, is that this is the only way to ensure a positive capital tax in the first period of the model – given the joint restrictions that are imposed on policy choice. The short-term benefits of taxing an inelastic factor are directly traded off against long-run costs. Under the Ramsey institutional design paradigm this trade-off can be avoided: a high initial capital tax and a low long-run rate can be selected independently. But this is precisely what gives Ramsey policy its unrealistic flavour: policymakers in practice do not legislate such directly time-contingent policy schemes into existence.

We have deliberately focused our attention in this paper on a model that is simple and already well understood, so as to extract the key implications of

[32] These assumptions are that the choice set should have a single degree of freedom everywhere, should allow – in principle – for any conceivable promise to be chosen in a given time period, and should ensure meaningful joint variation in promises at every horizon.

time-consistent institutional design as clearly as possible. But this has meant abstracting from one important problem: how to incorporate stochastic shocks into the analysis. We have shown that time-consistent institutional design in deterministic settings can only have at most one meaningful degree of freedom to vary policy promises in the different time periods. It is certainly not clear that this will remain the case when randomness affects the economy. Subsequent work will consider precisely this problem.

References

[1] Abreu, D., D. Pearce and E. Stacchetti (1990), 'Toward a Theory of Discounted Repeated Games with Imperfect Monitoring', *Econometrica*, 55(5), 1041–1063.

[2] Aiyagari, S.R. (1995), 'Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting', *Journal of Political Economy*, 103(6), 1158–1175.

[3] Alvarez, Y., J. Burbidge, T.Farrell and L. Palmer (1992), 'Optimal Taxation in a Life-Cycle Model', *Canadian Journal of Economics*, 25(1), 111–122.

[4] Atkeson, A., V.V. Chari, and P.J. Kehoe (1999), 'Taxing Capital Income: A Bad Idea', *Federal Reserve Bank of Minneapolis Quarterly Review*, 23(3), 3–17.

[5] Blake, A.P. (2001), 'A "Timeless Perspective" on Optimality in Forward-Looking Rational Expectations Models', NIESR Discussion Papers 188, National Institute of Economic and Social Research.

[6] Blake, A.P., and T. Kirsanova (2012), 'Discretionary Policy and Multiple Equilibria in LQ RE Models', *Review of Economic Studies*, 79(4), 1309–1339.

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

[8] Chari, V.V., and P.J. Kehoe (1999), 'Optimal Fiscal and Monetary Policy', In *Handbook of Macroeconomics Volume 1C*, ed. J.B. Taylor and M. Woodford, 1671–1745. Amsterdam: Elsevier Science, North-Holland.

[9] Conesa, J.C., S. Kitao and D. Krueger (2009), 'Taxing Capital Income: Not a Bad Idea After All!', *American Economic Review,* 99(1), 25–48.

[10] Debortoli, D., and R. Nunes (2010), 'Fiscal policy under loose commitment', *Journal of Economic Theory*, 145(3), 1005–1032.

[11] Debortoli, D., and R. Nunes (2013) ,'Lack Of Commitment And The Level Of Debt', *Journal of the European Economic Association*, 11(5), 1053–1078.

[12] Díaz-Giménez, J., G. Giovanetti, R. Marimon and P. Teles (2008), 'Nominal Debt as a Burden on Monetary Policy', *Review of Economic Dynamics*, 11(3), 493–514.

[13] Ellison, M., and N. Rankin, 'Optimal monetary policy when lump-sum taxes are unavailable: A reconsideration of the outcomes under commitment and discretion', *Journal of Economic Dynamics and Control*, 31(1), 219–243.

[14] Erosa, A., and M. Gervais (2002), 'Optimal Taxation in Life-Cycle Economies', *Journal of Economic Theory*, 105(2), 338–369.

[15] İmrohoroğlu, S. (1998), 'A Quantitative Analysis of Capital Income Taxation', *International Economic Review*, 39(2), 307–328.

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

[17] Klein, P., P. Krusell and J.-V. Ríos-Rull (2008), 'Time-Consistent Public Policy', *Review of Economic Studies*, 75(3), 789–808.

[18] Klein, P., and J.-V. Ríos-Rull (2003), 'Time-Consistent Optimal Fiscal Policy', *International Economic Review*, 44(4), 1217–1246.

[19] Kydland, F.E., and E.C. Prescott (1977), 'Rules Rather Than Discretion: The Inconsistency of Optimal Plans', *Journal of Political Economy*, 85(3), 473–491.

[20] Marcet, A., and R. Marimon (2014), 'Recursive Contracts', manuscript.

[21] Martin, F. (2009), 'A Positive Theory of Government Debt', *Review of Economic Dynamics*, 12(4), 608–631.

[22] Mele, A. (2011), 'Repeated Moral Hazard and Recursive Lagrangeans', MPRA Paper 30310, University Library of Munich, Germany.

[23] Niemann, S., P. Pichler and G. Sorger, 'Public Debt, Discretionary Policy, and Inflation Persistence', *Journal of Economic Dynamics and Control,* 37(6), 1097–1109.

[24] Ortigueira, S. (2006), 'Markov-Perfect Optimal Taxation', *Review of Economic Dynamics*, 9, 153–178.

[25] Reis, C. (2013), 'Taxation without Commitment', *Economic Theory*, 52(2), 565–588.

[26] Samuelson, P.A. (1958), 'An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money', *Journal of Political Economy*, 66(6), 467–482.

[27] Schaumburg, E., and A. Tambalotti (2007), 'An Investigation of the Gains from Commitment in Monetary Policy', *Journal of Monetary Economics*, 54(2), 302–324.

[28] Song, Z., K. Storesletten and F. Zilibotti (2012), 'Rotten Parents and Disciplined Children: A Politico-Economic Theory of Public Expenditure and Debt', *Econometrica*, 80(6), 2785–2803.

[29] Svensson, L E.O. (1999), 'How Should Monetary Policy Be Conducted in an Era of Price Stability?', in *New Challenges for Monetary Policy*, Kansas City: Federal Reserve Bank of Kansas City, 195–259.

[30] Woodford, M. (1999), 'Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?' in *New Challenges for Monetary Policy*, Kansas City: Federal Reserve Bank of Kansas City, 277–316.

[31] Woodford, M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.

[32] Woodford, M. (2010), 'Optimal Monetary Stabilization Policy', In *Handbook of Macroeconomics Volume 3*, ed. B.M. Friedman and M. Woodford, 723–828. Amsterdam: Elsevier Science, North-Holland.

[33] Yared, P. (2010), 'Politicians, Taxes, and Debt', Review of Economic Studies, 77(2), 806–840.

Appendix

7.1 Assumption on preferences

We place the following technical restriction on u. This is not central to our arguments, and will be used only to ensure the uniqueness of the time-consistent steady state:

Condition 1 *Consider any two allocations* (c', l') *and* (c'', l'')*, with* $(c'', l'') > (c', l')$ *(in the product order sense). Let* $w(\alpha)$ *be the marginal rate of substitution between consumption and leisure at an allocation that is a convex combination of these two:*

$$w(\alpha) := -\frac{u_l(\alpha c'' + (1-\alpha) c', \alpha l'' + (1-\alpha) l')}{u_c(\alpha c'' + (1-\alpha) c', \alpha l'' + (1-\alpha) l')} \tag{34}$$

We assume:

$$\frac{d^2 w(\alpha)}{d\alpha^2} > 0 \tag{35}$$

Proof of Proposition 2

Proof. Let $\theta \in \Theta \subset \mathbb{R}$ describe an arbitrary parameterisation of the promise values, such that $\{\omega_t(\theta)\}_{t=0}^{\infty} \in \Omega_0$ for all $\theta \in \Theta$, $\{\omega_t(\theta^*)\}_{t=0}^{\infty} = \widehat{\omega}^{0,\infty}$ for some $\theta^* \in \theta$, and $\omega_t(\theta)$ is continuously differentiable in θ when $\theta = \theta^*$. Let $\tilde{\theta}$ describe an alternative parameterisation with exactly the same properties, and $\tilde{\theta}^*$ defined correspondingly. We wish to show that if choice in Ω_0 satisfies time consistency then the differential changes $\left\{ \frac{d\omega_t(\theta)}{d\theta}\Big|_{\theta=\theta^*} \right\}_{t=0}^{\infty}$ cannot be linearly independent from their $\tilde{\theta}$ equivalents $\left\{ \frac{d\omega_t(\tilde{\theta})}{d\tilde{\theta}}\Big|_{\tilde{\theta}=\tilde{\theta}^*} \right\}_{t=0}^{\infty}$. Consider first the change induced by θ. Since $\widehat{\omega}^{0,\infty}$ is optimal in Ω_0, by the envelope theorem we must have:

$$\sum_{t=0}^{\infty} \beta^t \left[\beta\lambda_t^1 \left. \frac{d\omega_{t+1}(\theta)}{d\theta} \right|_{\theta=\theta^*} - \lambda_t^2 \left. \frac{d\omega_t(\theta)}{d\theta} \right|_{\theta=\theta^*} \right] = 0 \tag{36}$$

where λ_t^1 and λ_t^2 are the multipliers that follow from the day-to-day policy problem when $\widehat{\omega}^{0,\infty}$ is chosen. The assumption that $V_0^d\left(\widehat{\omega}^{0,\infty}, k_0\right) > -\infty$ implies there is a choice in Ω_0 that allows for a non-empty choice set in the day-to-day problem, and hence these multipliers are well defined. If the problem is non-trivial there must also exist at least one multiplier that is strictly positive. If not the day-to-day policymaker is unconstrained by the optimal institution in Ω_0, but if this is true then the day-to-day policymaker must select the first-best allocation subject to the resource constraint (19) alone. Our starting position is that this is not attainable.

By time consistency an equivalent condition to (36) must hold for variations in Ω_s for all $s \geq 0$:

$$\sum_{t=s}^{\infty} \beta^{t-s} \left[\beta\lambda_t^1 \left. \frac{d\omega_{t+1}(\theta)}{d\theta} \right|_{\theta=\theta^*} - \lambda_t^2 \left. \frac{d\omega_t(\theta)}{d\theta} \right|_{\theta=\theta^*} \right] = 0 \tag{37}$$

This is only possible if changes to θ induce zero marginal loss on a period-by-period basis:

$$\beta\lambda_t^1 \left. \frac{d\omega_{t+1}(\theta)}{d\theta} \right|_{\theta=\theta^*} - \lambda_t^2 \left. \frac{d\omega_t(\theta)}{d\theta} \right|_{\theta=\theta^*} = 0 \tag{38}$$

An equivalent condition is required for the changes induced by $\tilde{\theta}$:

$$\beta\lambda_t^1 \left. \frac{d\omega_{t+1}\left(\tilde{\theta}\right)}{d\tilde{\theta}} \right|_{\tilde{\theta}=\tilde{\theta}^*} - \lambda_t^2 \left. \frac{d\omega_t\left(\tilde{\theta}\right)}{d\tilde{\theta}} \right|_{\tilde{\theta}=\tilde{\theta}^*} = 0 \tag{39}$$

Since (38) and (39) are evaluated at the same point, the shadow values $\left\{\lambda_t^1, \lambda_t^2\right\}$ are identical across the two equations for all t. It follows that if they both hold for all t then either $\left\{ \frac{d\omega_t(\theta)}{d\theta}\Big|_{\theta=\theta^*} \right\}_{t=0}^{\infty}$ must equal $\left\{ \alpha \cdot \frac{d\omega_t(\tilde{\theta})}{d\tilde{\theta}}\Big|_{\tilde{\theta}=\tilde{\theta}^*} \right\}_{t=0}^{\infty}$, for some $\alpha \in \mathbb{R}$. Hence we cannot have linear independence. ■

Proof of Proposition 3

Proof. If allocations converge to a steady state then, from (21) and (22), the optimal promises in Ω_0 must do likewise, and so too the multipliers λ_{ss}^1, λ_{ss}^2 and μ_{ss}. It follows from assumption A1 that optimal choice within Ω_0 can be treated as the selection of a single parameter $\theta \in \Theta \subseteq \mathbb{R}$, where the promises are C^1 with respect to this parameter. By continuity and assumption A2 the optimal choice of θ cannot be at the boundary of Θ unless the optimal promises are of infinite magnitude, which is easily seen to be inconsistent with resource feasibility – and hence incompatible with steady state. It follows that a first-order condition will characterise the optimal choice of θ if any steady state is obtained, and by time consistency this condition must match (38) above:

$$\left.\frac{d\omega_{t+1}(\theta)}{d\theta}\right|_{\theta=\theta^*} = \frac{\lambda_{ss}^2}{\beta\lambda_{ss}^1}\left.\frac{d\omega_t(\theta)}{d\theta}\right|_{\theta=\theta^*} \tag{40}$$

where θ^* is the optimal parameter choice, and we allow for now the possibility that $\left.\frac{d\omega_t(\theta)}{d\theta}\right|_{\theta=\theta^*}$ need not be constant in t. But the continuity of $\frac{d\omega_t(\theta)}{d\theta}$ for all t means that assumption A3 would be violated unless $\frac{\lambda_{ss}^2}{\beta\lambda_{ss}^1} = 1$ (otherwise the effect of changing θ on ω_t relative to ω_0 would be growing or decaying without bound in t). Hence in any steady state $\left.\frac{d\omega_t(\theta)}{d\theta}\right|_{\theta=\theta^*}$ will be constant in t, and we must have:

$$\lambda_{ss}^2 = \beta\lambda_{ss}^1 \tag{41}$$

This condition, together with constraints (19), (21) and (22) and the first-order conditions from the day-to-day problem, (23) to (25), provide a set of seven non-linear equations in seven unknowns. Generically this system will have a finite number of solutions. We show below that Condition 1 is in fact sufficient to guarantee uniqueness in the solution, conditional on existence. But this part of the proof is quite involved, so for the sake of the general reader more interested in the positivity of steady-state capital taxes we demonstrate this claim first.

Substituting expression (41) into a steady-state version of the first-order condition for optimal choice of capital in the day-to-day problem, (25), gives:

$$\beta\left(1 + F_{k,ss} - \delta\right) = 1 + (1-\beta)\frac{\lambda_{ss}^1}{\mu_{ss}}u_{c,ss} \tag{42}$$

A necessary condition for optimal consumer choice in the steady state of a decentralised economy with capital tax τ_{ss}^k is:

$$\beta\left[1 + \left(1 - \tau_{ss}^k\right)\left(F_{k,ss} - \delta\right)\right] = 1 \tag{43}$$

Hence the optimum in the restricted choice set will have a positive capital tax iff:

$$(1-\beta)\frac{\lambda_{ss}^1}{\mu_{ss}}u_{c,ss} > 0 \tag{44}$$

The resource constraint clearly binds, so to guarantee positive capital taxes we need to rule out $\lambda^1_{ss} = 0$. Suppose this were true. Then the implementability condition would be slack:

$$\beta^{-1} u_{c,ss} k_{ss} < u_{c,ss} (c_{ss} + k_{ss}) + u_{l,ss} l_{ss} \tag{45}$$

which means that the steady-state allocation would be first best. In particular, we would have:

$$u_{l,ss} + u_{c,ss} F_{l,ss} = 0 \tag{46}$$

and:

$$\beta (1 + F_{k,ss} - \delta) = 1 \tag{47}$$

Substituting the previous two expressions into (45) gives:

$$(F_{k,ss} - \delta) k_{ss} < c_{ss} - F_{l,ss} l_{ss} \tag{48}$$

or, using the homogeneity properties of F:

$$c_{ss} > F (k_{ss}, l_{ss}) - \delta k_{ss} \tag{49}$$

This is clearly incompatible with resource feasibility.

It remains to show that the steady state is unique. This is difficult, ultimately because of the well-known problem in Ramsey tax settings that the implementability constraint need not describe a convex set of allocations, and this allows the possibility of local optima that are not globally best. We can nonetheless prove it under Condition 1 as follows. The set of seven equations characterising steady state can equivalently be obtained as a set of necessary requirements for an interior solution to the following 'static' problem:[33]

$$\max_{c,l,k} u (c, l)$$

subject to:

$$c + g + \delta k \leq F (k, l) + \frac{\beta - 1}{\beta} k \tag{50}$$

and:

$$u_c k \leq \beta [u_c (c + k) + u_l l] \tag{51}$$

where η is set to be the multiplier on (50) and λ^1 on (51), and λ^2 and ω can be defined by:

$$\lambda^2 = \beta \lambda^1 \tag{52}$$

and:

$$\omega = \beta^{-1} u_c k \tag{53}$$

respectively. Consider two distinct allocations $\{c', l', k'\}$ and $\{c'', l'', k''\}$ that are both feasible and implementable,[34] and deliver the same value for $u (c, l)$.

[33] We drop the steady-state subscripts for convenience here.
[34] That is, satisfy (50) and (51).

Without loss of generality, suppose $(c'', l'') > (c', l')$. Suppose we can show that any convex combination between $\{c', l', k'\}$ and $\{c'', l'', k''\}$ will also be feasible and implementable, for all such choices. Then it will follow by usual arguments that any first-order necessary conditions are also sufficient for the unique global optimum. Hence any solution to the seven equations that characterise steady state must also be unique.

The feasibility constraint (50) clearly describes a convex set of allocations, so if $\{c', l', k'\}$ and $\{c'', l'', k''\}$ are feasible then $\{(1-\alpha)c' + \alpha c'', (1-\alpha)l' + \alpha l'', (1-\alpha)k' + \alpha k''\}$ must be likewise, for all $\alpha \in [0,1]$. It remains to show that implementability is preserved at intermediate points. Without loss we focus on the case in which $\{c', l', k'\}$ and $\{c'', l'', k''\}$ both satisfy (51) with equality.[35] In this case we have:

$$c' - \left(\beta^{-1} - 1\right) k' + \frac{u_l'}{u_c'} l' = 0 \tag{54}$$

and:

$$c'' - \left(\beta^{-1} - 1\right) k'' + \frac{u_l''}{u_c''} l'' = 0 \tag{55}$$

For implementability to hold at intermediate points, we need:

$$[(1-\alpha)c' + \alpha c''] - \left(\beta^{-1} - 1\right)[(1-\alpha)k' + \alpha k''] - w(\alpha)[(1-\alpha)l' + \alpha l''] \geq 0 \tag{56}$$

where $w(\alpha)$ is as defined in the statement of Condition 1. Using the preceeding two expressions in this to eliminate the terms in c and k, and rearranging, it is sufficient that we have the following:

$$\alpha l'' [w(\alpha) - w''] + (1-\alpha) l' [w(\alpha) - w'] \leq 0 \tag{57}$$

Since $(c'', l'') > (c', l')$, and utility is the same for the two allocations, we must have $w'' > w(\alpha) > w'$ for all $\alpha \in (0,1)$, and $w(\alpha)$ is a monotone increasing function on $[0,1]$. By Condition 1 $w(\alpha)$ is additionally convex. Hence by Jensen's inequality we have:

$$\begin{aligned}
& \alpha l'' [w(\alpha) - w''] + (1-\alpha) l' [w(\alpha) - w'] & (58)\\
\leq\ & \alpha l'' [[(1-\alpha) w' + \alpha w''] - w''] + (1-\alpha) l' [[(1-\alpha) w' + \alpha w''] - w'] & (59)\\
=\ & -\alpha (1-\alpha)(l'' - l')(w'' - w') & (60)\\
<\ & 0 & (61)
\end{aligned}$$

as required. ■

7.2 Proof of Proposition 4

Proof. The proof works by showing that when assumptions A1 to A5 hold, it is impossible for $\widehat{\omega}_t(k_0)$ to depend on k_0. If $\widehat{\omega}^{0,\infty}(k_0)$ is a time-consistent optimal

[35] If the inequality were strict then a local movement along the line between the two points would clearly preserve strict implementability for values of α close enough to 0 or 1. If there were some values $\{\alpha', \alpha''\}$ at which equality held, the arguments in the main text can be applied relative to these points.

plan, we know from Proposition 3 that either $\widehat{\omega}_t$ must converge to ω^*_{ss} or real allocations do not converge. The latter case is ignored by the assumption of the proposition. Invariance and convergence to ω^*_{ss} are possible only if $\widehat{\omega}_t = \omega^*_{ss}$ for all t.

Taking the total derivative of $V_0^d\left(\widehat{\omega}^{0,\infty}(k_0), k_0\right)$ with respect to k_0, we have:

$$\begin{aligned}\frac{dV_0^d\left(\widehat{\omega}^{0,\infty}(k_0), k_0\right)}{dk_0} &= \frac{\partial V_0^d\left(\widehat{\omega}^{0,\infty}(k_0), k_0\right)}{\partial k_0} \\ &+ \sum_{t=0}^{\infty} \frac{\partial V_0^d\left(\widehat{\omega}^{0,\infty}(k_0), k_0\right)}{\partial \widehat{\omega}_t(k_0)} \frac{d\widehat{\omega}_t(k_0)}{dk_0}\end{aligned} \tag{62}$$

But from Assumption A5, switching to a promise series associated with an alternative initial capital stock cannot improve the policymaker's welfare. Hence we also have:

$$\sum_{t=0}^{\infty} \frac{\partial V_0^d\left(\widehat{\omega}^{0,\infty}(k_0), k_0\right)}{\partial \widehat{\omega}_t(k_0)} \frac{d\widehat{\omega}_t(k_0)}{dk_0} = 0 \tag{63}$$

or:

$$\sum_{t=0}^{\infty} \beta^t \left[\beta \lambda_t^1 \frac{d\widehat{\omega}_{t+1}(k_0)}{dk_0} - \lambda_t^2 \frac{d\widehat{\omega}_t(k_0)}{dk_0}\right] = 0 \tag{64}$$

By time invariance, a similar condition must hold for the promises observed from period s onwards:

$$\sum_{t=0}^{\infty} \beta^t \left[\beta \lambda_t^1 \frac{d\widehat{\omega}_{t+1}(k_s^*(k_0))}{dk_s^*(k_0)} - \lambda_t^2 \frac{d\widehat{\omega}_t(k_s^*(k_0))}{dk_s^*(k_0)}\right] = 0 \tag{65}$$

where $k_s^*(k_0)$ is the optimal capital stock for period t given $\widehat{\omega}^{0,\infty}(k_0)$. By assumption A4 we have:

$$\widehat{\omega}_t(k_s^*(k_0)) = \widehat{\omega}_{t+s}(k_0) \tag{66}$$

And thus:

$$\frac{d\widehat{\omega}_t(k_s^*(k_0))}{dk_s^*(k_0)} \frac{dk_s^*(k_0)}{dk_0} = \frac{d\widehat{\omega}_{t+s}(k_0)}{dk_0} \tag{67}$$

Setting $s = 1$, pre-multiplying the left-hand side of (65) by $\beta \frac{dk_s^*(k_0)}{dk_0}$ and using the result in (64) gives:

$$\beta \lambda_0^1 \frac{d\widehat{\omega}_1(k_0)}{dk_0} = \lambda_0^2 \frac{d\widehat{\omega}_0(k_0)}{dk_0} \tag{68}$$

The time invariance assumption implies, by symmetry, that for all t we must also then have:

$$\beta \lambda_t^1 \frac{d\widehat{\omega}_1(k_t^*(k_0))}{dk_t^*(k_0)} = \lambda_t^2 \frac{d\widehat{\omega}_0(k_t^*(k_0))}{dk_t^*(k_0)} \tag{69}$$

so:

$$\beta\lambda_t^1 \frac{d\widehat{\omega}_{t+1}\left(k_t^*\left(k_0\right)\right)}{dk_t^*\left(k_0\right)} \frac{dk_t^*\left(k_0\right)}{dk_0} = \lambda_t^2 \frac{d\widehat{\omega}_t\left(k_t^*\left(k_0\right)\right)}{dk_t^*\left(k_0\right)} \frac{dk_t^*\left(k_0\right)}{dk_0} \tag{70}$$

$$\beta\lambda_t^1 \frac{d\widehat{\omega}_{t+1}\left(k_0\right)}{dk_0} = \lambda_t^2 \frac{d\widehat{\omega}_t\left(k_0\right)}{dk_0} \tag{71}$$

From this it follows that either $\frac{d\widehat{\omega}_t(k_0)}{dk_0} = 0$ for all t, or $\frac{d\widehat{\omega}_0(k_0)}{dk_0}$ takes a non-zero value, and for all $t \geq 0$ $\frac{d\widehat{\omega}_{t+1}(k_0)}{dk_0}$ satisfies the recursion:

$$\frac{d\widehat{\omega}_{t+1}\left(k_0\right)}{dk_0} = \frac{\lambda_t^2}{\beta\lambda_t^1} \frac{d\widehat{\omega}_t\left(k_0\right)}{dk_0} \tag{72}$$

Suppose this latter condition is indeed true. Convergence implies:

$$\lim_{t\to\infty} \left[\frac{d\widehat{\omega}_t\left(k_0\right)}{dk_0}\right] = 0 \tag{73}$$

That is, the sequence induced by (72) converges to zero. But notice that an identical recursion is satisfied by $\frac{d\widehat{\omega}_t(k_0)}{d\theta}$: the derivative of $\widehat{\omega}_t$ along the unique degree of freedom available to the policymaker. We have:

$$\frac{d\widehat{\omega}_{t+1}\left(k_0\right)}{d\theta} = \frac{\lambda_t^2}{\beta\lambda_t^1} \frac{d\widehat{\omega}_t\left(k_0\right)}{d\theta} \tag{74}$$

for all $t \geq 0$ and arbitrary non-zero $\frac{d\widehat{\omega}_0(k_0)}{d\theta}$. But by assumption A3 this sequence cannot converge to zero. This is a contradiction, leaving $\frac{d\widehat{\omega}_t(k_0)}{dk_0} = 0$ as the only possibility. ∎