Time-Consistent Institutional Design*

Charles Brendon† Martin Ellison‡

January 11, 2015

This paper reconsiders normative policy design in environments subject to time inconsistency problems, à la Kydland and Prescott (1977). In many such settings, the time-varying dynamics of Ramsey policy make it unrealistic as a practical policy option. An alternative approach is needed to deliver useful policy recommendations. We propose a choice criterion that is more general than optimality, but is 'time-consistent' in the sense that it can be applied recursively. It is a version of the Pareto principle, taken with respect to the preference orderings of policymakers at different points in time. The idea is that commitment to an 'institution' should not be regretted by all current and future policymakers. The criterion implies choices that differ from steady-state outcomes under Ramsey policy. For example, we recommend modest positive steady-state capital taxes in a problem similar to Judd (1985), where steady-state Ramsey capital taxes are zero.

Keywords: Institutional Design; Pareto Efficiency; Ramsey Policy; Time-Consistency

JEL Codes: D02, E61

*This is a very substantially revised version of a manuscript previously 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.

†Faculty of Economics and Queens' College, University of Cambridge. Email: cfb46@cam.ac.uk

‡Department of Economics and Nuffield College, University of Oxford. Email: martin.ellison@economics.ox.ac.uk

1. Introduction

Time-inconsistency problems of the type first highlighted by Kydland and Prescott (1977) are a dominant feature of the modern macroeconomic policy literature. Whenever expectations of future policy influence the feasible set of current choices, there will be an incentive *ex-ante* for a policymaker to issue promises that it is not *ex-post* optimal to keep. This presents a dilemma for policy design. No dynamic path exists that is recursively optimal, i.e., best at the start of time and best when viewed as a subsequent continuation outcome.

Since recursive optimality is impossible, a distinction emerges between positive and normative policy analysis, based on the timing of choice. Positive analysis considers the consequences of period-by-period decisionmaking by successive generations of policymakers, each taking as given the actions of their successors. Outcomes are generically sub-optimal, with welfare losses particularly large when non-Markov reputational equilibria are ruled out.[1] Normative analysis instead considers how to design policy when choice can be made once-and-for-all at the start of time. It assumes that there is a commitment device to prevent any reneging on past promises.

The current paper is normative in this sense. Unlike the vast majority of the normative literature, however, its focus is not on policies that are 'Ramsey-optimal', i.e. best from the perspective of the initial time period. Our main contribution is instead to present an alternative choice criterion, which, we argue, can deliver policies that are more appealing for practical purposes than the Ramsey solution.

Kydland and Prescott problems, by definition, are such that no dynamic policy choice can ever be best for all time periods. In response, the conventional Ramsey approach selects a policy that is best for one particular period – the first.[2] This is a straightforward way to resolve differences in policy preferences over time, but it comes with significant practical drawbacks. Because choice is tailored to be best for the initial period only, the resulting policies often feature implausible dynamics, or long-run outcomes that seem far from desirable, or both. Our approach to the problem is different. Since no policy exists that is best for all periods, we ask whether there exist choices that satisfy a weaker criterion than optimality in all periods. This alternative criterion is a version of the Pareto principle, taken with respect to the preferences of policymakers at different points in time. We find that there do, and that these policies have appealing properties relative to the Ramsey solution.

Our general motivation is that in many settings the Ramsey solution does not seem to give a useable answer to the question: 'How should policy be designed?' Numerous examples exist. When

[1] On reputational equilibria see, in particular, Chari and Kehoe (1990), Atkeson (1991), Phelan and Stachetti (2001) and Golosov and Iovino (2014). The influential paper by Klein, Krusell and Ríos-Rull (2008) refers to Markov-perfect outcomes as 'time-consistent policy'. Our focus will be on normative policies that exhibit time consistency, so we avoid this useage. Markov-perfect equilibria have been the focus of a large recent literature, including papers by Klein and Ríos-Rull (2003), Ortigueira (2006), Ellison and Rankin (2007), Díaz-Giménez et al. (2008), Martin (2009), Reis (2013), Niemann et al. (2013) and Blake and Kirsanova (2012).

[2] As is well known, Ramsey policy can be computed by recursive techniques once the state vector is augmented to include either direct promise values or shadow costs associated with past constraints. Kocherlakota (1996) provided an early application of the former, whilst the latter has been pioneered by Marcet and Marimon (2014) who show that the Lagrange multipliers in incentive problems are equivalent to Pareto weights. These techniques do not imply that Ramsey policy is recursive in the sense of being time-consistent, as there is always an incentive *ex-post* to ensure additional promise-keeping constraints do not bind.

designing capital taxes in the style of Chamley (1986) and Judd (1985), for instance, Ramsey policy generally involves punitively high wealth taxes initially, converging to zero as time progresses.[3] In dynamic models of social insurance under asymmetric infomation, significant initial degrees of redistribution are often matched with the long-run immiseration of almost all agents under Ramsey policy.[4]

No matter how realistic the underlying models, neither of these proposals seems likely to be accepted by a real-world policymaker. The problem is not simply that policy changes over time, but that it does so independently of the evolution of the natural state vector. There is an inherent asymmetry in the way time periods themselves are treated when designing Ramsey policy. Choice is best for the initial time period, 'period zero', only. This is passed through as an asymmetry in the outcomes of policy. For the same natural state vector, choice in the initial period will be different from choice at a later date.

Many economists have expressed unease about this way of structuring choice. Svensson (1999) puts the problem succinctly, asking 'Why is period zero special?' Long-term legislation relating to tax policy, social insurance structures, monetary policy and other macroeconomic instruments is commonplace in all major economies. What is not common is for this legislation to impart dynamics to policy that are not intrinsic to the economy itself. Perhaps this is because policymakers find it impossible to commit to desirable policy plans even when writing long-lasting legislation. Another alternative, however, is that lawmakers place independent value on a choice procedure that does not treat any one time period as 'special'. If this is so, policy advice will be redundant unless allowances are made for it. By retaining recursive applicability, the purpose of our choice criterion is to do just this.

To date the literature has addressed problematic aspects of Ramsey policy on a largely case-by-case basis, depending on whether it is the long-run outcome or transition dynamics that appear more implausible. The immiseration result, for instance, has motivated Phelan (2006) and Farhi and Werning (2007, 2010) to attach distinct non-zero Pareto weights to later generations when designing social insurance policy. This effectively amounts to an increase in the social discount factor above its private-sector value.[5] As an approach this is indeed sufficient to overturn immiseration, but its implications stretch far wider. In particular, it implies changing policy in models in which no Kydland and Prescott problem is present, such as a textbook Ramsey growth model. A long-running debate, ranging from Ramsey himself (1928) to Stern (2006), considers the normative appropriateness of discounting future generations' welfare. But this is a far broader problem than time inconsistency à la

[3]Straub and Werning (2014) have recently shown that in some versions of the Judd (1985) model a Ramsey-optimal plan may converge to a 'corner' steady-state with zero consumption, rather than zero capital taxes. Below we consider simulations in a representative-agent model with variable labour supply, in which case capital taxes are zero in the long run for all conventional assumptions on preference parameters. In either case the dynamic asymmetry in policy is apparent.

[4]Thomas and Worrall (1990) were first to highlight the immiseration result, in the context of a moral hazard model. Kocherlakota (2010) provides a useful discussion on the requirements for immiseration to obtain in the context of a dynamic Mirrlees (hidden type) model.

[5]Related work by Sleet and Yeltekin (2006) has shown that the best sustainable outcome in an environment when policymakers are completely myopic coincides with Ramsey-optimal policy when the policymaker is more patient than the private sector. This provides a link between positive and normative approaches, though the motivation for Phelan (2006) and Farhi and Werning (2007, 2010) is principally normative.

Kydland and Prescott, and it would be preferable to keep it separate.

In the New Keynesian monetary policy literature, by contrast, it is the fact that Ramsey policy comes with transition dynamics that is identified as problematic. To address this, Woodford (1999, 2003) has advocated a 'timeless' approach to policy design, which involves implementing steady-state Ramsey policy from the start of time.[6] The capital tax literature often proceeds in similar fashion, albeit more informally, with high taxes along the transition generally being neglected for the purposes of policy advice.[7] But the justification for such an approach unclear. There is no particular reason why the steady state of Ramsey policy should itself be desirable, independently of the transition. What if the long-run outcome of Ramsey policy is immiseration for almost all agents?

The approach in this paper is different from both of these literatures. Unlike Phelan (2006) and Farhi and Werning (2007, 2010), we consider Pareto efficiency exclusively with respect to the time-inconsistent aspects of policy choice. Without a Kydland and Prescott problem, the choice procedure that we propose will yield standard solutions. Unlike the New Keynesian and capital tax literatures, our method is specifically designed to deliver policy that is 'transition-free', rather than simply discarding the transition dynamics from the Ramsey plan *ex-post*.

Central to our analysis is a novel decomposition of Kydland and Prescott problems into two distinct components. The first component, which we interpret as day-to-day policymaking, is a choice problem across conventional policy instruments, taking as given a basic set of targets for policy, or 'promises'.[8] This choice problem is entirely time-consistent. The second component is a choice problem over the promises themselves. This can be interpreted as an institutional design problem.

The institutional design problem is always time-inconsistent in the standard sense. That is, no sequence of promises is recursively optimal. Yet this does not rule out the recursive application of a weaker criterion than optimality. This is the approach we take in order to arrive at a time-consistent institutional design procedure. Instead of optimal, in each period we require that institutional choice should be Pareto efficient with respect to the preferences of all current and future generations of policymakers. Given the commitment that is being implemented, there should not exist an alternative set of promises that every generation of policymakers would prefer to switch to. To take a simple example, suppose an inflation target had been set in perpetuity at two per cent per annum, but every generation would like the ability to switch to a permanent target of four per cent instead. Then our Pareto criterion would fail.

In general this criterion will be satisfied by a large number of policies in any given period, including Ramsey-optimal policy from that period on. The important point is that there will now exist policies that are contained in the expanded choice set in *every* period. We call such policies recursively Pareto efficient. Unlike Ramsey policy, recursively Pareto efficient policies need only depend on the natural state vector of the economy.

A general feature of policies that satisfy recursive Pareto efficiency is that they differ systematically from Ramsey policy, including in steady state. This is an important result, as it suggests that the

[6]More recent papers in the New Keynesian tradition that follow this approach include Adam and Woodford (2012), Benigno and Woodford (2012), and Corsetti, Dedola and Leduc (2010). Damjanovic, Damjanovic and Nolan (2008) offer an alternative approach based on maximising steady-state welfare.

[7]See, for instance, the influential survey paper by Atkeson, Chari and Kehoe (1999).

[8]For example, a central bank determines day-to-day monetary policy to best fulfill its institutional mandate.

steady-state aspects of Ramsey policy, such as zero capital tax rates, cannot be justified as 'desirable' independently of the policy transition. The main reason for the difference is that Ramsey policy places too great a weight on past promises in steady state, relative to our Pareto criterion.[9]

We define and characterise recursively Pareto efficient policies both in a general setting and by reference to three well-known examples. The first is a simple linear-quadratic New Keynesian inflation bias problem, the second a capital tax problem in the style of Judd (1985)[10], and the third a model of social insurance with participation constraints, in the spirit of Kocherlakota (1996). In the inflation bias example, policy satisfying recursive Pareto efficiency delivers the lowest loss in the set of time-invariant policies. This is superior to implementing the steady-state Ramsey policy.[11] With the capital tax example, Ramsey policy delivers capital taxes that start close to 300 per cent of net income, before gradually decaying to zero.[12] This occurs even when the capital stock starts at its long-run steady-state level. Under recursively Pareto efficient policy, tax rates differ from their steady-state levels only if the capital stock is also away from steady state. If there are transition dynamics in the capital stock, the associated departure of tax rates from steady state is small by comparison with the Ramsey case. Steady-state capital tax rates are around 20 to 30 per cent of net income for conventional parameterisations.

The social insurance example is one in which a fixed fraction of the population is assumed to receive a high-income draw with a given probability each period, whereas the remainder of the population receives permanently low income. Ramsey policy for a utilitarian planner sees the effective Pareto weight of current high earners increased each period, a stationary increase that ultimately implies there is no redistribution in steady state to the part of the population that has permanently low income. The recursive Pareto efficiency criterion instead has the Pareto weights of past high earners drifting downwards, which guarantees that there will be redistribution to permanently low earners in steady state.

2. General setup

Time is discrete, and runs from period 0 to infinity.[13] In each period t there exists a policymaker with preferences over allocations from period t onwards. These allocations are of the form $\{x_{s+1}, a_s\}_{s=t}^{\infty}$ where $x_s \in X \subset \mathbb{R}^n$ is a vector of n states determined in period $s-1$ and $a_s \in A \subset \mathbb{R}^m \times \mathbb{R}^{\varsigma}$ is a vector of controls determined in period s. There are m controls, and each is defined for all possible realisations of a stochastic vector $\sigma_s \in \Sigma$, where Σ is a countable set of cardinality ς, which may be infinite. $a_s(\sigma_s) \in \mathbb{R}^m$ denotes the value of a_s particular to the realisation σ_s of the stochastic process. The value of x_t is given at the start of period t. The policymaker's preferences are described by a

[9]More technically, Marcet and Marimon (2014) have shown that the accumulated shadow values associated with keeping past promises are generally non-stationarity under Ramsey policy. When policy is recursively Pareto efficient this is no longer true.

[10]That is, an example in which the government must abide by a balanced-budget constraint each period.

[11]A related observation was made by Blake (2001), when considering the stochastic steady-state implied by Ramsey policy in the dynamic New Keynesian model with cost-push shocks.

[12]It is assumed that taxes can be levied in excess of net income, i.e., that the underlying asset itself can be taxed.

[13]Notation is adapted from Marcet and Marimon (2014).

time-separable objective criterion W_t:

$$W_t := \sum_{s=t}^{\infty} \beta^{s-t} r\left(x_s, a_s\right). \tag{1}$$

The policymaker is constrained by a set of n restrictions defining the evolution of the state vector:

$$x_{s+1} = l\left(x_s, a_s\right), \tag{2}$$

a set of i contemporaneous restrictions linking controls and states:

$$p\left(x_s, a_s\right) \geq 0, \tag{3}$$

and 'forward-looking' constraints of two different types. The first is a set of j time-separable restrictions looking forward over the infinite horizon:

$$\mathbb{E}_s \sum_{\tau=1}^{\infty} \beta^{\tau} h\left(a_{s+\tau}\left(\sigma_{s+\tau}\right), \sigma_{s+\tau}\right) + h_0\left(a_s\left(\sigma_s\right), \sigma_s\right) \geq 0, \tag{4}$$

whereas the second is a set of k restrictions looking forward one period ahead:

$$\mathbb{E}_s \beta g_1\left(a_{s+1}\left(\sigma_{s+1}\right), \sigma_{s+1}\right) + g_0\left(a_s\left(\sigma_s\right), \sigma_s\right) \geq 0. \tag{5}$$

It is assumed that $j + k \geq 1$ so the policymaker is subject to at least one forward-looking constraint, but otherwise there is no requirement that any of i, j, k or m should be non-zero. Constraints (2) to (5) hold for all $s \geq t$ in period t, with the functions in the constraints vector-valued and of the specified dimension. The values of the functions in (4) and (5) are allowed to vary directly in the vector of the exogenous stochastic process σ_s, as well as indirectly through $a_s\left(\sigma_s\right)$. Where the meaning is clear, we will usually keep the dependence of a_s on σ_s implicit, writing $h\left(a_s, \sigma_s\right)$ and so on.

The expectations in (4) and (5) are taken with respect to an ergodic Markov process for σ_s defined on Σ. The process has a time-invariant probability of transiting from state σ to state σ' that is denoted by $P\left(\sigma'|\sigma\right)$ and a stationary probability of state σ denoted by $P\left(\sigma\right)$. Constraints (4) and (5) must hold for all initial $\sigma_s \in \Sigma$. By assumption there is no aggregate uncertainty. This approach to modelling uncertainty is somewhat restrictive, but is economical on notation and sufficient to incorporate a number of important models. This includes the social insurance example below, where there is idiosyncratic income risk. It will be useful at times to make reference to $A\left(\sigma\right) \subset \mathbb{R}^m$ as the space of control variables available for a given σ.

The distinction between restrictions that are infinite-horizon and one-period ahead in (4) and (5) is made because these are the two main forms of forward-looking constraints in most problems of interest. This is without loss of generality. Intermediate cases in which outcomes in s are constrained by expectations of outcomes up to finite horizon $s + n$ can be written in the form of (5), by defining appropriate auxiliary variables and additional contemporaneous restrictions of the form in (2) linking these auxiliary variables to the controls. Similarly, the absence of state variables from forward-looking constraints is without loss of generality when appropriate auxiliary variables and additional contemporaneous restrictions are likewise defined.

It is constraints of the form (4) and (5) that are responsible for time inconsistency. The policymaker in period t is not restricted to ensure that the versions of these constraints relating to $t-1$ and earlier remain satisfied, and in general it will be best to renege on any past promises to do so. This was the problem highlighted by Kydland and Prescott (1977). The next section briefly introduces three examples that can be nested in the general setup.

2.1. Three examples

Example 1: A linear-quadratic inflation bias problem

Allocations in period t consist of inflation π_t and the output gap y_t. These are both control variables so there are no state variables. The policymaker's objective is:

$$-\sum_{s=t}^{\infty} \beta^{s-t} \left[\pi_s^2 + \chi \left(y_s - \bar{y} \right)^2 \right], \tag{6}$$

where $\chi > 0$ is a parameter and $\bar{y} > 0$ is the optimal level for the output gap.[14] The policymaker is subject to a single constraint, a deterministic version of the New Keynesian Phillips curve. This is a one-period ahead restriction of the form in (5):

$$\pi_s = \beta \pi_{s+1} + \gamma y_s, \tag{7}$$

with γ a parameter. This forward-looking constraint provides an incentive to promise low inflation in the future so as to ease the current inflation-output trade-off. Such a promise is generally time-inconsistent. Full details can be found in Woodford (2003).

Example 2: A capital tax problem

This is a variant of the balanced-budget problem studied by Judd (1985). A representative agent consumes, saves and supplies labour. For exogenous reasons, the government must consume a fixed quantity g of real resources each period. Government consumption is funded by a linear tax on labour income and a linear tax on capital income net of depreciation. The government cannot borrow. Allocations in period t are consumption c_t, labour l_t, output y_t and the capital stock k_t. The capital stock is the only state variable and the rest are controls. The policymaker's objective in period t is to maximise the lifetime utility of the representative agent:

$$\sum_{s=t}^{\infty} \beta^{s-t} u \left(c_s, l_s \right) \tag{8}$$

[14] This exceeds the natural level of output due to monopoly power in the product market.

The policymaker faces an aggregate resource constraint of the form in (2):

$$k_{s+1} = y_s - c_s - g + (1 - \delta) k_s, \tag{9}$$

and a production constraint of the form in (3):

$$y_s \leq F(k_s, l_s). \tag{10}$$

The distortionary character of taxes implies a further 'implementability' constraint that restricts allocations available to the policymaker when the problem is written as here in its primal form.[15] In the present case this is a one-period ahead forward-looking restriction of the form in (5):[16]

$$\beta \{u_{c,s+1}(c_{s+1} + k_{s+2}) + u_{l,s+1} l_{s+1}\} \geq u_{c,s} k_{s+1} \tag{11}$$

This constraint says that the value of consumption and capital purchases in period $t + 1$, net of any labour income that period, must weakly exceed the value of the capital holdings that are taken into period $t + 1$, where these values are calculated in period t at prices corresponding to anticipated marginal rates of substitution. In short, it prevents the policymaker from restricting the consumer's spending power *ex post* relative to what is anticipated. This will conflict with the incentive of a later policymaker to tax the consumer's existing, inelastic capital holdings.

Example 3: Social insurance with participation constraints

This is a variant of the limited commitment model due to Kocherlakota (1996). There is a continuum of agents indexed on the unit interval, with each agent receiving an endowment each period. Measure $\mu \in [0, 1)$ of agents receive a low income y^l in every period. The remaining measure $(1 - \mu)$ receive a high income $y^h > y^l$ with probability p in a given period, and a low income y^l with probability $(1 - p)$. The endowment draws are independent across agents and time, and publicly observable. The Ramsey-optimal plan in this environment has the consumption levels of agents subject to income risk depending only on the time elapsed since those agents last received a high-income draw, so attention is restricted to policies with this feature.[17] The exogenous stochastic variable $\sigma_{s,i} \in \Sigma$ can then be defined as the number of periods since agent i last drew a high income, with Σ the set of positive integers (including 0). The Markov process governing $\sigma_{s,i}$ for generic agent i is:

$$\sigma_{s,i+1} = \begin{cases} \sigma_{s,i} + 1 & \textit{with prob } (1 - p) \\ 0 & \textit{with prob } p \end{cases}.$$

[15] See Chari and Kehoe (1999).

[16] Strictly speaking, this differs from (5) in containing a variable dated at $s + 2$. However, said variable can be eliminated using the resource constraint (9).

[17] This is done mainly to keep the notation simple, since for these policies it is sufficient for the policymaker to summarise an agent's history up to period s in the single variable $\sigma_{s,i}$. The results obtained do not change when more general forms of dependence are formally allowed.

Agents subject to income risk are only differentiated by the time since they last had a high-income draw, and hence can be indexed by their $\sigma_{s,i}$ values.[18] Furthermore, the stochastic process governing $\sigma_{s,i}$ implies that there will be measure $(1-\mu)(1-p)^{\sigma}p$ of agents in each period who last received a high-income draw σ periods ago. The policymaker's objective in period t is utilitarian:

$$\sum_{s=t}^{\infty}\beta^{s-t}\left[(1-\mu)\sum_{\sigma=0}^{\infty}(1-p)^{\sigma}pu\left(c_s\left(\sigma\right)\right)+\mu u(c_s^p)\right], \tag{12}$$

where $c_s(\sigma)$ is the consumption in period s of an agent who received a high-income draw σ periods ago, and c_s^p is the consumption in period s of an agent who has a permanently low income. The utility function u satisfies the usual properties. The problem has interesting properties without incorporating public or private asset accumulation, so the resource constraint is assumed to hold period by period:

$$(1-\mu)\sum_{\sigma=0}^{\infty}(1-p)^{\sigma}pc_s(\sigma)+\mu c_s^p\leq[1-(1-\mu)p]y^l+(1-\mu)py^h. \tag{13}$$

Incentive compatibility requires the insurance scheme to deliver at least as much utility to an agent as they could obtain in autarky, which implies infinite horizon forward-looking constraints of the form in (4):

$$\mathbb{E}_t\sum_{s=t}^{\infty}\beta^{s-t}u\left(c_s\left(\sigma_{s,i}\right)\right)\geq u\left(y_t\right)+\frac{\beta}{1-\beta}\left[pu(y^h)+(1-p)u(y^l)\right], \tag{14}$$

for agents drawing income y_t, and

$$\sum_{s=t}^{\infty}\beta^{s-t}u(c_t^p)\geq\frac{1}{1-\beta}u(y^l) \tag{15}$$

for the permanently low income agents. These incentive compatibility constraints are the source of time inconsistency. The government has an *ex ante* incentive to minimise costs and maintain the social insurance scheme by promising high future utility to agents with a high income draw, but there will be *ex post* incentives to renege on these promises.

2.2. Assumptions

It is useful to describe a range of possible restrictions that can be placed on general functions r, l, p, h, h_0, g_1 and g_0 later in the analysis.

Assumption 1. *The functions $r : X\times A\to\mathbb{R}$, $l : X\times A\to\mathbb{R}^n$, $p : X\times A\to\mathbb{R}^i$, $h : A(\sigma)\to\mathbb{R}^j$, $h_0 : A(\sigma)\to\mathbb{R}^j$, $g_1 : A(\sigma)\to\mathbb{R}^k$ and $g_0 : A(\sigma)\to\mathbb{R}^k$ are continuous. The spaces $A\subset\mathbb{R}^m$ and $X\subset\mathbb{R}^n$ are compact and convex.*

Assumption 2. *The functions r, l, p, h, h_0, g_1 and g_0 are continuously differentiable.*

[18]For simplicity, it is assumed that information on the infinite history of endowment draws is available from the start of time. This information will be irrelevant to the Ramsey plan, but may be of use in designing a stationary policy. It would make no difference to the argument if this 'information' were fictitiously drawn according to the true underlying distribution in period 0.

Assumption 3. *The functions h, h_0, g_1 and g_0 are quasi-concave.*

Assumption 4. *The function r is strictly concave, the function p is quasi-concave and the function l is linear.*

Assumption 5. *The functions p, h, h_0, g_1 and g_0 are concave.*

Assumption 1 brings some basic structure to the problem and will be assumed throughout. In most environments of interest the relevant constraint functions are utility functions, production functions, profit functions and the like, for which continuity is a relatively innocuous imposition. Convexity and compactness are similarly conventional structures to impose on the spaces A and X. Assumption 2 is invoked principally for ease of exposition. It would be possible to relax it, but only at notational cost and without changing the character of the results, so it is likewise imposed throughout. Assumptions 3, 4 and 5 are progressively stronger. They are needed for instance in the capital tax example, where it is unclear that the relevant functions in the implementability constraint (11) satisfy quasi-concavity, let alone full concavity. Nonetheless, simple transformations of variables can sometimes allow a problem that appears to violate Assumption 3 to be written in a form that satisfies even Assumption 5. This is the case in the capital tax example if the utility function is isoelastic and separable.

The main purpose of Assumptions 3, 4 and 5 is to allow some structure to be placed on a class of indirect objective functions to be defined below. These will be central to the analysis of alternative normative strategies, and differing assumptions on the primitives will affect what can be said about efficient resolutions to the time inconsistency problem.

2.3. Ramsey policy

Ramsey policy is an allocation $\left\{x_{s+1}^R, a_s^R\right\}_{s=0}^{\infty}$ that maximises W_0 subject to all relevant constraints of the form (2) to (5) for all $s \geq 0$. In general, the continuation of this policy $\left\{x_{s+1}^R, a_s^R\right\}_{s=t}^{\infty}$ from period $t > 0$ will not maximise W_t subject to the same constraints being satisfied for $s \geq t$. This is because the Ramsey policy will have been influenced by a desire to affect expectational constraints that were binding in periods prior to t, but which are no longer a concern – the time inconsistency problem. The characteristics of Ramsey policy in the three examples are next introduced and discussed.

Example 1: A linear-quadratic inflation bias problem

The Ramsey plan for the linear-quadratic inflation bias example is familiar from the New Keynesian literature.[19] Figure 1 shows the dynamic paths of inflation and output under a conventional calibration of $\beta = 0.96$, $\gamma = 0.024$, $\chi = 0.048$, $\bar{y} = 0.1$. Initial choices are unconstrained by the effects of current inflation on past expectations, meaning that the costs of engineering high output are initially quite low. The output gap is initially set above 9 per cent, after which it is optimal to allow inflation to drift downwards over time as lower future inflation permits higher current output under the New Keynesian Phillips Curve (7). In steady state the inflation rate is zero, as is the output gap.

[19]See, for example, Woodford (2003).

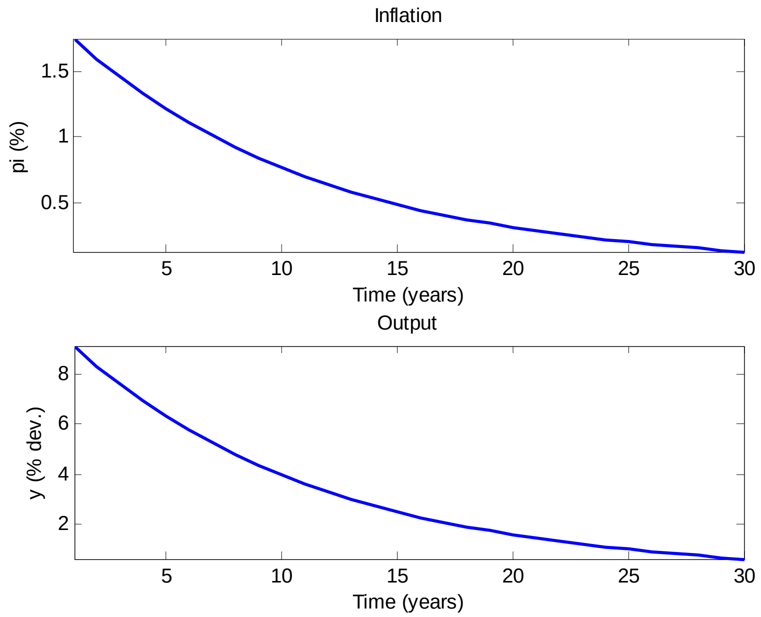

Figure 1: Ramsey paths for inflation and output in the linear-quadratic inflation bias problem

Example 2: A capital tax problem

The Ramsey path in the capital tax example is plotted in Figure 2, assuming standard parameter values and functional forms.[20] The initial capital stock is set equal to the Ramsey steady-state value, and plotted in the lower panel as a percentage of the steady-state capital stock. The broad properties of the Ramsey path are familiar from the literature following Chamley (1986) and Judd (1985). Initial taxes on net capital income are implausibly high, at around 300 per cent, but decay quickly and converge on zero as time progresses.[21] The capital stock reflects the path of capital taxes, falling over the first 10 years to a level about 5 per cent below its steady-state value as higher taxes reduce the incentives to save. The capital stock only gradually recovers afterwards as taxes fall and the incentives to save are restored.

Example 3: Social insurance with participation constraints

The general properties of social insurance models with participation constraints have been explored in a number of recent papers.[22] Figure 3 charts the dynamic consumption path of individuals whose income endowment is constrained to be low each period.[23] These agents initially receive a significant

[20] Specifically, utility takes an additively separable isoelastic form. Consumption utility is logarithmic and labour disutility is exponential, with an inverse Frisch elasticity equal to 2. The production function is Cobb-Douglas with capital share 0.33. $\beta = 0.96$, $\delta = 0.05$ and $g = 0.6$, which ensures a steady-state government consumption to output ratio of 0.31.

[21] There is no economic reason to rule out capital taxes in excess of 100 per cent, as agents can always meet the associated liabilities by selling their underlying capital holdings.

[22] Among others, see Krueger and Perri (2006), Krueger and Uhlig (2006), Broer (2013) and Ábrahám and Laczó (2014).

[23] The calibration uses log consumption utility with $\beta = 0.96$, $\mu = 0.2$, $p = 0.01$, $y^l = 1$ and $y^h = 10$. Qualitative outcomes are not strongly dependent on these choices.

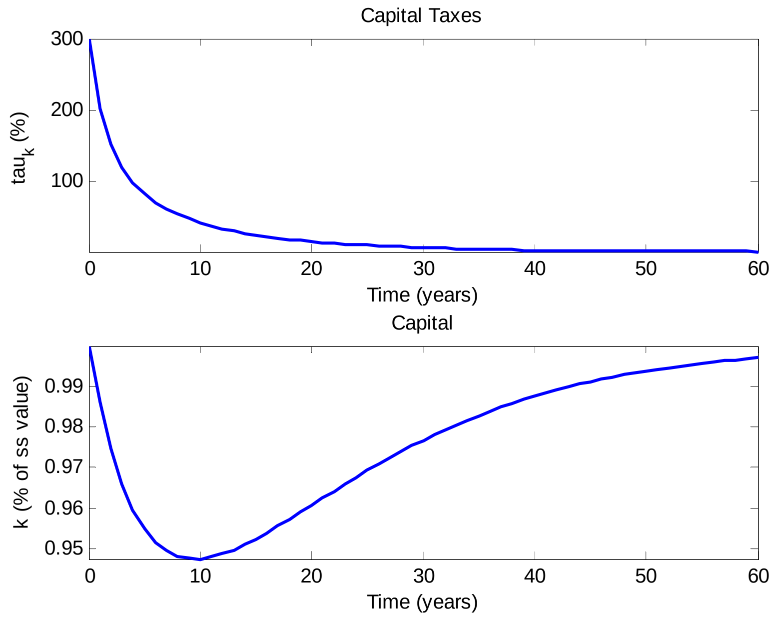

Figure 2: Ramsey capital taxes and the capital stock in the capital tax problem

transfer, raising their consumption to just below the average endowment in the economy.[24] As time progresses, their consumption drifts down as transfers are instead directed towards satisfying the participation constraints of agents who have just received a high-income draw. The permanently low income agents are eventually limited to consuming only their endowment. An identical consumption trajectory is followed by an agent who is subject to income risk but is unlucky enough always to draw a low income.

Discussion: Asymmetric objectives and outcomes

An important characteristic of Ramsey policy in all three examples is its dynamic asymmetry. Ramsey-optimal inflation trends downwards in the linear-quadratic inflation bias example, even though the structure of this simple New Keynesian economy is entirely stationary. The same is true of optimal consumption for permanently low income agents in the social insurance example. The capital tax example features capital as an endogenous state variable, which could potentially account for some of the dynamics in policy choices. However, the calibration assumes that the initial capital stock is equal to its eventual steady-state value. Nonetheless, the character of the initial policy with capital taxes at 300 per cent could hardly be more different from that which obtains in the limit when capital taxes converge on zero. Thus in all three examples the Ramsey policy induces a different allocation in identical economic circumstances, dependent entirely on the amount of time that has progressed since optimisation took place.

The dynamic asymmetry of outcomes is unsurprising given the definition of Ramsey policy. No policy exists that is optimal in the set of possibilities from every time period onwards. The Ramsey

[24] The average endowment is 1.072. A first-best utilitarian policy would provide this level of consumption to all agents in all periods.

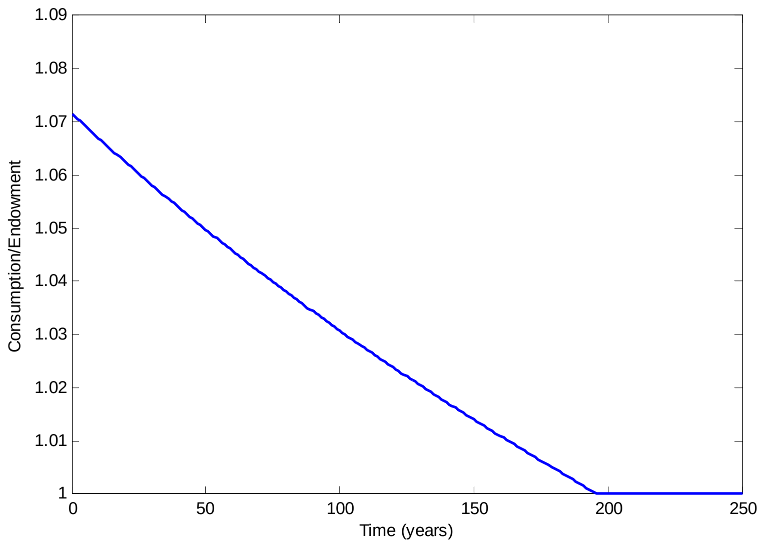

Figure 3: Consumption path for low-income agents in the social insurance problem

approach selects a policy that is best from period 0 onwards, but this policy cannot also be best from generic period s onwards, otherwise there would be no time-inconsistency problem. This implies that if the natural state of the economy is identical in periods 0 and s then policies must differ between the two periods. Such asymmetry is an uncomfortable property for policy models to exhibit. It is difficult, for instance, to imagine that a central bank might pursue a time-varying inflation target, or that capital taxes would follow the sort of trajectory exhibited in Figure 2.

What is needed is an approach to designing policy that delivers time-symmetric allocations whilst remaining meaningfully optimal. The question is what this concept ought to be. The 'timeless perspective' approach of Woodford (1999, 2003) proposes one answer that has been widely applied in the New Keynesian policy literature.[25] It requires the policymaker in period 0 to immediately implement the steady-state allocations associated with Ramsey policy. Thus, Woodford (2003) would advocate an optimal inflation rate of zero in period 0 of the linear-quadratic inflation bias example. His heuristic justification is that the Ramsey steady-state policy is time-invariant, and would have formed part of a Ramsey-optimal path had optimisation taken place in the distant past. The resulting choices would have been optimal from *some* time perspective, albeit one prior even to the first period of the model. The same focus on Ramsey steady-state outcomes has been emphasised more informally in the dynamic capital tax literature, where policy recommendations have typically stressed the zero steady-state capital tax rates whilst discarding the associated transition to steady state.[26]

However, 'Do what would have been planned for today in the distant past' may not always be a desirable, or even feasible, maxim to follow. The example of social insurance with participation constraints sounds a particularly cautionary note. The permanently low income agents receive no

[25]Recent examples include Adam and Woodford (2012) and Corsetti, Dedola and Leduc (2010). Woodford (2010) provides a more detailed discussion of the merits of eliminating asymmetries over time.

[26]See Atkeson et al. (1999). Straub and Werning (2014) have recently highlighted the inseparability of transition dynamics from the optimality of the subsequent zero rate.

transfers from their more fortunate peers in Ramsey steady state, and so are left consuming their low income endowment forever. Immediate implementation of the steady-state allocations would hence eliminate all the gains from redistribution that accrue during the first 200 'early' years of the Ramsey policy. This is clearly not a desirable outcome. It is similarly unclear whether the Ramsey steady-state allocation can necessarily be implemented in period 0 when the policy environment includes endogenous state variables. In a version of the social insurance problem with public storage and a sufficiently high real interest rate, the Ramsey-optimal policy involves the policymaker accumulating sufficient assets so that first-best complete risk-sharing is achieved in the long run.[27] This clearly cannot be implemented from the very first time period if the policymaker has not yet accumulated sufficient assets. In examples such as these the timeless perspective policy does not seem well defined.[28]

The approach we take is to define a policy criterion that is weaker than 'optimality' but can still be applied recursively. We argue that a version of the Pareto criterion can do just this, providing a new systematic approach to institutional and policy design that delivers desirable time-symmetric allocations. To make progress, the next section distills the time inconsistency problem by separating out the underlying choice variables to which time inconsistency is attached.

3. Inner and outer problems

The general problem can be divided into two components, an 'inner' and 'outer' problem.[29] The outer problem is concerned with the selection of a dynamic path for promise variables. These correspond to the promise values used by Abreu, Pearce and Stachetti (1990) to give a recursive structure to the Ramsey problem, though they are not used here directly to obtain recursivity. Instead, the outer problem is one of choosing an entire time-contingent sequence of such promises, which we label an institutional design problem. An 'inner' problem can then be cast, determining the optimal dynamic allocation subject to the promise constraints given by the outer problem. The value of this inner problem is contingent on the given sequence of promises, and can be interpreted as an indirect utility function across possible promise sequences. The time inconsistency problem then manifests itself as time variation in this preference structure, and conventional normative criteria, such as Pareto efficiency, can be used to assess the desirability of alternative institutions.

[27] Ljungqvist and Sargent (2012), Chapter 20, gives a textbook presentation of this result.

[28] From a technical perspective, it is known from the work of Marcet and Marimon (2014) that the Pareto weights of agents are non-stationary in the social insurance example with participation constraints. Woodford (2003) presents his approach as imposing steady-state values for the multipliers on past promises when solving the period 0 decision problem. But these multipliers are one and the same as the Pareto weights, and thus cannot take steady-state values because of their non-stationarity.

[29] The separation is analogous to that of Hansen and Sargent (2007) in their specification of robust control theory.

3.1. The inner problem

The inner problem solves for optimal allocations, conditional on feasibility and given sequences for the promise values. Mathematically, in period $t \geq 0$ the inner problem solves:

$$\max_{\{x_{s+1}, a_s\}_{s=t}^{\infty}} W_t := \sum_{s=t}^{\infty} \beta^{s-t} r\left(x_s, a_s\right),$$

subject to the evolution of the state vector (2), the restrictions linking controls and states (3), and the promise constraints:

$$\begin{aligned} h_0\left(a_s, \sigma_s\right) + \beta \mathbb{E} \omega_{s+1}^h\left(\sigma_{s+1}\right) &\geq 0, && (16)\\ h\left(a_s, \sigma_s\right) + \beta \mathbb{E} \omega_{s+1}^h\left(\sigma_{s+1}\right) &\geq \omega_s^h\left(\sigma_s\right), && (17)\\ g_0\left(a_s, \sigma_s\right) + \beta \mathbb{E} \omega_{s+1}^g\left(\sigma_{s+1}\right) &\geq 0, && (18)\\ g_1\left(a_s, \sigma_s\right) &\geq \omega_s^g\left(\sigma_s\right), && (19) \end{aligned}$$

for all $s \geq t$, where $x_t \in X$ is the initial state vector and $\omega_s^h\left(\sigma_s\right) \in \mathbb{R}^j$ and $\omega_s^g\left(\sigma_s\right) \in \mathbb{R}^k$ are sequences of σ-contingent vectors for the promise values in all $s \geq t$.

The promise values ω_s^h and ω_s^g are stacked in the vector $[\omega_s^h\left(\sigma_s\right)', \omega_s^g\left(\sigma_s\right)']' := \omega_s\left(\sigma_s\right) \in \mathbb{R}^{j+k}$ and the collection of promise values across σ_s draws is denoted by $\omega_s := \left\{\omega_s\left(\sigma_s\right)\right\}_{\sigma_s \in \Sigma}$. Conditions (16) and (18) are 'promise-making' constraints, as they restrict the choice of variables in the h_0 and g_0 functions relative to a vector of future promise values. The implication is that future values of the h and g_1 functions must be at least as large as these promise values, in order for the full expectational constraints to be satisfied. This is ensured by the 'promise-keeping' constraints (17) and (19).

For the original constraints (4) and (5) to hold for all $s \geq t$, it is sufficient to ensure that (16) and (18) hold for all $s \geq t$ and (17) and (19) for all $s > t$. Imposing constraints (17) and (19) also for period t would add an additional restriction on choice. This can be motivated by a need to respect past promises, but not by the fundamental economic restrictions that feature from period t onwards.

3.1.1. Notation

The set of infinite promise sequences $\{\omega_s\}_{s=t}^{\infty}$ such that the inner problem has a non-empty constraint set is denoted by $\Omega\left(x_t\right) \subset \left(\mathbb{R}^{j+k} \times \mathbb{R}^{\varsigma}\right)^{\infty}$, for an initial state vector x_t. The interior of $\Omega\left(x_t\right)$ is represented by $\mathring{\Omega}\left(x_t\right)$. Compactness of A and X and the continuity properties imposed under Assumption 1 together imply that the constraint functions in (16) to (19) are bounded uniformly in s. This in turn means that sufficiently large, uniform bounds can be imposed on the values of $\omega_s^h\left(\sigma_s\right)$ and $\omega_s^g\left(\sigma_s\right)$ without affecting the problem. These bounds are incorporated into the definition of $\Omega\left(x_t\right)$:[30]

Assumption. $\{\omega_s\}_{s=t}^{\infty} \in \Omega\left(x_t\right)$ *only if the sequence* $\{\omega_s\}_{s=t}^{\infty}$ *is uniformly bounded in* s *for all* $x_t \in X$.

The value of the inner problem is denoted by $V(\{\omega_s\}_{s=t}^{\infty}, x_t)$, for all $x_t \in X$ and all $\{\omega_s\}_{s=t}^{\infty} \in \Omega\left(x_t\right)$. This object will be a major focus of the analysis, and we label it the 'promise-value function'. Where

[30]The condition is stated as an assumption for clarity. It is left unnumbered since it is without loss of additional generality.

convenient, the convention will be that $V(\{\omega_s\}_{s=t}^{\infty}, x_t) = -\infty$ when the inner problem has an empty constraint set, which allows V to be defined on the entire product space $\left(\mathbb{R}^{j+k} \times \mathbb{R}^{\varsigma}\right)^{\infty}$. It will also be useful to consider the evolution of the state variables associated with any given promise sequence. We say that the sequence $\{\omega_s\}_{s=t}^{\infty}$ and initial state x_t 'induces' a sequence of state vectors $\{x'_{s+1}\}_{s=t}^{\infty}$ if there exists a sequence of control variables $\{a'_s\}_{s=t}^{\infty}$ such that $\{x'_{s+1}, a'_s\}_{s=t}^{\infty}$ solves the inner problem for given $\{\omega_s\}_{s=t}^{\infty}$ and x_t.

3.1.2. Time consistency

Proposition 1. *The inner problem is **time consistent**. That is, if $\{x'_{s+1}, a'_s\}_{s=t}^{\infty}$ solves the inner problem for promise sequence $\{\omega_s\}_{s=t}^{\infty}$ and initial state vector x_t, then the continuation $\{x'_{s+1}, a'_s\}_{s=t+\tau}^{\infty}$ solves the inner problem for promise sequence $\{\omega_s\}_{s=t+\tau}^{\infty}$ and initial state vector $x'_{t+\tau}$ for all $\tau \geq 1$.*

The proof is straightforward and given in the appendix. Time consistency is an important property of the inner problem. The essential point is that time inconsistency problems derive from different policymakers having different incentives to make, keep and renege upon *promises*. Once these promises are treated as given, no additional source of time inconsistency remains.

3.2. Properties of the promise-value function

The promise-value function $V(\{\omega_s\}_{s=t}^{\infty}, x_t)$ is not an object commonly analysed in the literature. It is new and plays a central role in the arguments that follow, so it is prudent to briefly consider some of its properties under varying combinations of Assumptions 1 to 5 defined earlier. Proofs rely on established properties of parameterised optimisation problems, and are contained in the appendix.

Proposition 2. *Suppose Assumptions 1 and 3 hold. Fix x_t and $\{\omega_s\}_{s=t}^{\infty} \in \mathring{\Omega}(x_t)$. The promise-value function $V(\cdot, x_t)$ is **continuous** at $\{\omega_s\}_{s=t}^{\infty}$.*[31]

The proof of continuity relies on a standard application of Berge's Theory of the Maximum. Continuity is an important regularity property for the V function to satisfy, but it will be helpful later to strengthen it to continuous differentiability. This requires a standard constraint qualification condition to be satisfied by the restrictions (16) to (19) at a chosen allocation, a linear independence constraint qualification (LICQ) that is presented in the appendix[32]. It amounts to requiring that each binding constraint in (16) to (19) is affected in a linearly independent manner by changes in the policy variables. This is not a significant limitation in any of the applications we study.

Proposition 3. *Suppose Assumptions 1 to 4 h:old. Fix x_t and suppose that LICQ is satisfied at the solution to the inner problem for all $\{\omega_s\}_{s=t}^{\infty} \in \mathring{\Omega}(x_t)$ and x_t. Then the promise-value function $V(\{\omega_s\}_{s=t}^{\infty}, x_t)$ is **continuously differentiable** in each element of the sequence $\{\omega_s\}_{s=t}^{\infty}$. Its derivative with respect to $\omega_s(\sigma_s)$ is given by the $(j+k) \times 1$ vector:*

$$-\beta^{s-t} P(\sigma_s) \begin{bmatrix} \lambda_s^h(\sigma_s) \\ \lambda_s^{g,1}(\sigma_s) \end{bmatrix} + \beta^{s-t} P(\sigma_s) \begin{bmatrix} \lambda_{s-1}^{h,0}(\sigma_{s-1}) + \lambda_{s-1}^{h}(\sigma_{s-1}) \\ \lambda_{s-1}^{g,0}(\sigma_{s-1}) \end{bmatrix} \tag{20}$$

[31]Continuity is defined with respect to the sup-norm, see proof.
[32]See, for instance, Wachsmuth (2013).

where $\lambda_s^{h,0}(\sigma_s)$, $\lambda_s^{h}(\sigma_s)$, $\lambda_s^{g,0}(\sigma_s)$ and $\lambda_s^{g,1}(\sigma_s)$ are the vector multipliers on constraints (16) to (19) respectively, and:

$$\lambda_{t-1}^{h,0}(\sigma_{t-1}) = \lambda_{t-1}^{h}(\sigma_{t-1}) = \lambda_{t-1}^{g,0}(\sigma_{t-1}) = 0$$

The main contribution of Proposition 3 is to establish the conditions under which a standard envelope condition applies to the promise-value function V. In the event that it does, the derivatives associated with changes to the promise vectors are simple linear combinations of the multipliers on the promise-keeping and promise-making constraints.

The expressions for the derivatives in (20) directly reflect the time inconsistency associated with choice of the promises in the outer problem. So long as $s > t$, the second term will generically be non-zero because there are shadow marginal benefits from making a promise that will later be kept. When $s = t$ these benefits have passed and the marginal effect of changing contemporaneous promises is only a cost, the first term in (20) associated with the promise-keeping constraints. It is also worth noting that the marginal effect associated with the promise-making constraint in period $s-1$ is discounted at the same rate β^{s-t} as the marginal effect associated with the promise-keeping constraint in period s. This is due to the presence of β pre-multiplying future promises in constraints (16) to (18). It implies that multipliers will generally evolve in a non-stationary fashion when promises are chosen optimally from the perspective of period t, which implies that the derivative is set to zero. This non-stationarity property has been highlighted by the work of Marcet and Marimon (2014). It has important implications for the character of the Ramsey solution in the long run, particularly in models with dynamic incentive constraints such as the social insurance example.

Proposition 4. *Suppose Assumptions 1, 4 and 5 hold. Fix x_t. The promise-value function $V(\cdot, x_t)$ is* ***strictly quasi-concave*** *in $\{\omega_s\}_{s=t}^{\infty} \in \Omega(x_t)$. If Assumptions 1, 4 and 5 hold but r is only weakly concave then $V(\cdot, x_t)$ is* ***quasi-concave****.*

Proposition 5. *Suppose Assumptions 1, 4 and 5 hold. Fix x_t. The space $\Omega(x_t)$ is* ***convex****.*

The proofs of these two Propositions are near-identical, with the exception that Proposition 5 relates solely to the constraint set, so goes through without any concavity restrictions on r. For this reason we prove 4 only in the appendix. Quasi-concavity implies that upper contour sets in the space $\Omega(x_t)$ are convex, given that $\Omega(x_t)$ is likewise. This is of substantial use when establishing the Pareto ranking of alternative sequences of promises, for reasons familiar from textbook general equilibrium analysis.

3.3. The outer problem: Ramsey policy and time inconsistency

The outer problem is to choose a sequence for the promise values $\{\omega_s\}_{s=0}^{\infty} \in \Omega(x_0)$. Choice here is subject to a time inconsistency problem, because from the perspective of period t it will never be desirable for the promise vector ω_t to place a meaningful constraint on choice in the associated inner problem from period t onwards. Any benefits from issuing promises accrue in the time periods when the promises are made, not when they are kept. But the advantage of having separated the inner and outer problems is that time inconsistency can now be viewed as a dynamic inconsistency in policymakers' preference orderings over promise sequences. For a given state vector x_t the promise-value function $V\left(\{\omega_s\}_{s=t}^{\infty}, x_t\right)$ describes a rational preference ordering over the space $\left(\mathbb{R}^{j+k} \times \mathbb{R}^{\varsigma}\right)^{\infty}$. We use variation in these preference structures to provide a formal definition of time inconsistency:

Definition 1. Fix x_0, and consider a promise sequence $\{\omega_s'\}_{s=0}^{\infty} \in \Omega(x_0)$ that induces $\{x_{s+1}'\}_{s=0}^{\infty}$. We say that this promise sequence is **time-consistent** if and only if there exists no other sequence $\{\omega_s''\}_{s=0}^{\infty} \in \left(\mathbb{R}^{j+k} \times \mathbb{R}^{\varsigma}\right)^{\infty}$ such that $V(\{\omega_s''\}_{s=t}^{\infty}, x_t') > V(\{\omega_s'\}_{s=t}^{\infty}, x_t')$ for some $t \geq 0$.

A Ramsey promise sequence can be defined using the initial-period promise-value function:

Definition 2. Fix x_0. The promise sequence $\{\omega_s^R\}_{s=0}^{\infty} \in \Omega(x_0)$ comprises a **Ramsey plan** if and only if there exists no alternative sequence $\{\omega'_s\}_{s=0}^{\infty} \in \left(\mathbb{R}^{j+k} \times \mathbb{R}^{\varsigma}\right)^{\infty}$ such that $V(\{\omega_s'\}_{s=0}^{\infty}, x_0) > V(\{\omega_s^R\}_{s=0}^{\infty}, x_0)$.

The next proposition establishes what was first shown by Kydland and Prescott (1977):

Proposition 6. *Let $\{\omega_s^R\}_{s=0}^{\infty}$ be a Ramsey plan, given some x_0. Exactly one of the following is true:*

1. *Constraints (16) to (19) never bind in the inner problem, given $\{\omega_s^R\}_{s=0}^{\infty}$ and x_0.*
2. *The Ramsey plan is not time-consistent.*

This well-understood result does not require much further comment. Either promises never matter, or else keeping them is not a time-consistent choice. It is, though, instructive to characterise the Ramsey plan in terms of the derivatives of the V function. Assuming that the boundary of $\Omega(x_0)$ does not constrain choice, Ramsey policy solves the unconstrained problem of maximising $V(\{\omega_s\}_{s=t}^{\infty}, x_0)$ with respect to each element of the promise sequence. Provided the necessary conditions for differentiability are met, applying the results of Proposition 3 means a necessary optimality condition with respect to the choice of $\omega_s(\sigma_s)$ is:

$$\left[\begin{array}{c} \lambda_s^h(\sigma_s) \\ \lambda_s^{g,1}(\sigma_s) \end{array}\right] = \left[\begin{array}{c} \lambda_{s-1}^{h,0}(\sigma_{s-1}) + \lambda_{s-1}^{h}(\sigma_{s-1}) - \\ \lambda_{s-1}^{g,0}(\sigma_{s-1}) \end{array}\right], \tag{21}$$

for $s > 0$, and:

$$\left[\begin{array}{c} \lambda_0^h(\sigma_0) \\ \lambda_0^{g,1}(\sigma_0) \end{array}\right] = 0 \tag{22}$$

These results replicate the common finding that dynamic multipliers on expectational constraints generally exhibit non-stationarity. In models with participation constraints such as the social insurance example, this is equivalent to the set of cross-sectional Pareto weights applied across agents being non-decreasing over time. This observation is central to the recursive multiplier formulation of Ramsey policy due to Marcet and Marimon (2014). It implies that agents who receive a series of consecutive low income draws see their share of total resources diminish, exactly as in our social insurance example. Long-run outcomes may be particularly adverse for these individuals. It is likely that the allocation of resources in any steady state will be driven by the need to make good on past promises, rather than the maximisation of an underlying social welfare objective.

4. The Pareto approach to designing promises

The purpose of this section is to propose an alternative to the Ramsey approach for designing dynamic promise sequences. The benefit of expressing time inconsistency through the promise-value function is that $V(\{\omega_s\}_{s=t}^{\infty}, x_t)$ can be analysed as a standard preference ordering over promise sequences from period t onwards. This means that well-established normative criteria can be invoked to resolve differences in preferences, in particular the Pareto criterion. This section considers the implications of such a Pareto approach.

4.1. Pareto efficiency: alternative definitions

Ex-post versus *ex-ante* criteria

The promise-value functions $V(\{\omega_s\}_{s=t}^{\infty}, x_t)$ for different values of t provide alternative rankings over continuation promise sequences, given an inherited state vector. There are two ways that Pareto efficiency could be defined with respect to these functions, depending on the treatment of the state vector.[33] Applying an *ex-ante* criterion, the comparisons between promise sequences are under the assumption that each promise sequence induces its own sequence for the evolution of the state vector. With an *ex-post* criterion, one instead assumes that the evolution of the state vector is the same under both promise sequences up to and including the period at which the pairwise comparisons are being made. The *ex-ante* criterion implies the following definition of (strict) Pareto efficiency:

Definition. A promise sequence $\{\omega_s^*\}_{s=t}^{\infty}$ inducing state vector $\{x_{s+1}^*\}_{s=t}^{\infty}$ is ***ex-ante* strictly Pareto efficient** in period t, if there is no alternative promise sequence $\{\omega_s'\}_{s=t}^{\infty}$ inducing state vector $\{x_{s+1}'\}_{s=t}^{\infty}$ such that for all $\tau \geq 0$:

$$V(\{\omega_s'\}_{s=t+\tau}^{\infty}, x_{t+\tau}') \geq V(\{\omega_s^*\}_{s=t+\tau}^{\infty}, x_{t+\tau}^*),$$

with the inequality strict for at least one τ.[34]

Intuitively, the *ex ante* criterion assesses promise sequences as if they were an efficient outcome of a hypothetical bargaining process, conducted at the start of period t between all policymakers from period t onwards. The '*ex-ante*' terminology refers to the idea that the two sequences are implicitly being compared at the start of period t, prior to the realisation of the associated dynamic allocation.

The alternative criterion is *ex-post* Pareto efficiency:

Definition. A promise sequence $\{\omega_s^*\}_{s=t}^{\infty}$ inducing state vector $\{x_{s+1}^*\}_{s=t}^{\infty}$ is ***ex-post* strictly Pareto efficient** in period t, if there is no alternative promise sequence $\{\omega_s'\}_{s=t}^{\infty}$ such that for all $\tau \geq 0$:

$$V(\{\omega_s'\}_{s=t+\tau}^{\infty}, x_{t+\tau}^*) \geq V(\{\omega_s^*\}_{s=t+\tau}^{\infty}, x_{t+\tau}^*),$$

with the inequality strict for at least one τ.

[33] There are similarities here with the difficulty of applying the standard Pareto criterion in dynamic models with endogenous population. See, for instance, Golosov, Jones and Tertilt (2007).

[34] The value of the state vector in period t satisfies $x_t' = x_t^* = x_t$.

The difference compared to *ex-ante* efficiency is that the path for the state vector is now fixed at $\{x^*_{s+1}\}_{s=t}^{t+\tau}$ when making the comparison between $\{\omega^*_s\}_{s=t+\tau}^{\infty}$ and $\{\omega'_s\}_{s=t+\tau}^{\infty}$ in period $t+\tau$. *Ex-post* Pareto efficiency therefore provides a test of the 'regret' associated with a given commitment. When it fails, it means that every policymaker at every point in time would like to switch to a fixed alternative path for promises. This desire to switch is assessed after the state variables have evolved up to period $t+\tau$, according to the original commitment.

That every commitment will be regretted in time inconsistency problems is a trivial statement. For every t, the policymaker in period t would prefer to replace any inherited commitment $\{\omega^*_s\}_{s=t}^{\infty}$ with the sequence of promises for the Ramsey plan from that period onwards, $\{\omega^{R,t}_s\}_{s=t}^{\infty}$. However, this latter sequence is itself particular to the policymaker in period t. It is far from trivial to say there is a *fixed* alternative promise sequence $\{\omega'_s\}_{s=t}^{\infty}$ such that a switch from $\{\omega^*_s\}_{s=t+\tau}^{\infty}$ to $\{\omega'_s\}_{s=t+\tau}^{\infty}$ is preferred by the policymakers for every period $t+\tau$. When this is true, there is a 'uniformity' to regret. All policymakers regret the commitment relative to the alternative fixed promise sequence. This is different from requiring that a commitment should not be regretted in each $t+\tau$ relative to a distinct promise sequence that is chosen freely in each $t+\tau$. We will occasionally use the terminology 'uniform regret' to capture a situation in which *ex-post* Pareto efficiency fails.

Reasons to analyse *ex-post* Pareto efficiency

The focus for the remainder of the paper is on *ex-post* Pareto efficiency. The first reason is that the ex-post criterion is more closely connected to the notion of time consistency we are interested in. Recall that a promise sequence is time-consistent if there is no alternative sequence of promises, such that at least *some* policymaker at some time $t+\tau$ would prefer to switch to it. *Ex-post* Pareto efficiency instead holds if there is no alternative sequence of promises such that *every* policymaker from period t onwards would prefer to switch. In this regard, the *ex-post* Pareto efficiency criterion can be viewed as a generalisation of time consistency. A time-consistent promise sequence, if possible, would be *ex-post* Pareto efficient, but not *vice-versa*. Importantly, it is a property of rankings under precisely the same preference structures $V(\cdot, x^*_{t+\tau})$ that themselves exhibit time inconsistency. This is not true of the *ex-ante* criterion.

The second reason to focus on *ex-post* Pareto efficiency is that *ex-ante* notions have difficulty separating out issues of time inconsistency from the far broader question of how best to give weights to different generations. Farhi and Werning (2007 and 2010) consider the set of Pareto efficient allocations between generations in dynamic private information settings, based on a similar *ex-ante* Pareto criterion.[35] They show that the Pareto frontier can be described by increasing the social discount factor in an otherwise standard planning problem.[36] This certainly affects the manner in which dy-

[35] This criterion differs from the definition of *ex-ante* Pareto efficiency given above, as it relates to direct choice over allocations rather than over promise sequences (which in turn induce allocations). Indeed, if one did prefer an *ex-ante* notion of Pareto efficiency then there would be little sense in separating out the inner and outer problems in the first place. Viewed *ex-ante*, the accumulation of state variables in the inner problem is just as much a source of disagreement between policymakers in different periods as the choice of promises.

[36] Specifically, this traces out the frontier implied when the Pareto weights on all generations are positive, but decay geometrically over time. The effective social discount factor is then constant after the first period, and greater than the private-sector discount factor β.

namic asymmetric information problems are resolved. Farhi and Werning (2007) demonstrate that it is sufficient to overcome the immiseration result in an environment similar to Atkeson and Lucas (1992). However, it also has strong implications in environments where the only dynamic linkages come from standard endogenous state variables, such as a textbook Ramsey growth problem. Here, a higher social discount factor will generically imply more capital accumulation. The Farhi-Werning approach thus succeeds in overturning uncomfortable long-run outcomes associated with Ramsey policy under time inconsistency, but only at the cost of accepting unconventional solutions to problems where no time inconsistency is present. We find that the *ex-post* criterion does not have this disadvantage.

The remainder of the paper refers to *ex-post* Pareto efficiency as 'Pareto efficiency' for simplicity, unless there is a clear risk of ambiguity.

Weak versus strong criteria

An additional, more nuanced distinction is between strict and weak notions of Pareto efficiency, which will be important in what follows. The definitions above relate to the strict criterion. The definition of the weak *ex-post* criterion is:

Definition. A promise sequence $\{\omega_s^*\}_{s=t}^{\infty}$ inducing state vector $\{x_{s+1}^*\}_{s=t}^{\infty}$ is ***ex-post* weakly Pareto efficient** in period t, if there is no alternative promise sequence $\{\omega_s'\}_{s=t}^{\infty}$ such that for some $\varepsilon > 0$ and all $\tau \geq 0$:

$$V(\{\omega_s'\}_{s=t+\tau}^{\infty}, x_{t+\tau}^*) - V(\{\omega_s^*\}_{s=t+\tau}^{\infty}, x_{t+\tau}^*) \geq \varepsilon.$$

In words, weak Pareto efficiency requires only that there is no alternative promise sequence that would strictly improve outcomes for all policymakers. It is thus a refinement of the strict Pareto efficiency requirement that there should be no alternative that would be weakly better for all.[37]

4.2. Recursive Pareto efficiency

The Ramsey-optimality criterion cannot be applied recursively in environments with time inconsistency. This is the reason for our focus on Pareto efficiency, which does turn out to permit a recursive application. The definition of recursive Pareto efficiency is:

Definition. A promise sequence $\{\omega_s^*\}_{s=0}^{\infty}$ is **recursively strictly (weakly) Pareto efficient** (RSPE / RWPE) iff the continuation sequence $\{\omega_s^*\}_{s=t}^{\infty}$ is strictly (weakly) Pareto efficient for all $t \geq 0$.

Heuristically, recursive Pareto efficiency requires not just that a dynamic promise sequence should not be uniformly regretted when it is first implemented, but also that it should not *come* to be uniformly regretted later. There is no guarantee that Ramsey policy will satisfy either the strict or weak definition of recursive Pareto efficiency. It is Pareto efficient in period 0, but its continuation may not be thereafter.

[37] Writing this condition in terms of ε, rather than as a strict inequality, rules out the possibility that the two value functions converge to one another at the limit as $\tau \to \infty$. The definition therefore requires that there are no alternatives that are strict improvements both in finite time and at the limit. This broadens the definition of the Pareto set.

5. Recursively Pareto-efficient policies

This section explores the implications for policy of adopting the recursive criterion.

5.1. Strict Pareto efficiency: a negative result

The first result is a negative one: the strict Pareto criterion cannot be applied recursively. In general, it is impossible to find a policy that satisfies recursive strict Pareto efficiency. Formally:

Proposition 7. *Suppose the Ramsey plan is not time-consistent when the initial state vector is x_0. Then no promise sequence in $\Omega(x_0)$ satisfies RSPE.*

The proof is in the appendix. The intuition is straightforward because it in essence restates the time-inconsistency problem. Suppose that the policymaker in period t was required to adhere to a sequence of promises $\{\omega_s^*\}_{s=t}^{\infty}$ such that some elements of ω_t^* were binding. Then switching to a plan $\{\omega_t', \{\omega_s^*\}_{s=t+1}^{\infty}\}$ such that no elements of ω_t' were binding would be strictly preferred in period t. This switch would make the policymaker in period t better off, whilst leaving all future policymakers indifferent because the two continuation promise sequences coincide. Hence there is a violation of strict Pareto efficiency. It follows either that a desirable allocation can be achieved without any promise constraints ever binding – as when the Ramsey plan is time-consistent – or that strict Pareto efficiency cannot be applied recursively.

Proposition 7 demonstrates that the strict Pareto efficiency criterion cannot be implemented recursively, just as was the case with Ramsey optimality. However the proof is very particular to the strict definition of Pareto efficiency, which fails if the first policymaker can be made better off while all future policymakers are left indifferent. If the purpose of the Pareto criterion is to capture the idea that 'uniform regret' should be avoided, then the lack of a persistent improvement is problematic: indifference is a very weak form of regret. A weak Pareto criterion, if satisfied, would instead ensure that every policymaker did not strictly regret the chosen commitment.

5.2. Recursive weak Pareto efficiency: a steady-state result

The second result is a positive one: the weak Pareto criterion can be applied recursively. When it is applied, policy in steady state is characterised by the multipliers on the promise constraints (16) to (19). Formally:

Proposition 8. *Consider a promise sequence $\{\omega_s\}_{s=0}^{\infty}$ inducing a sequence of state vectors $\{x_{s+1}\}_{s=0}^{\infty}$ from initial state vector x_0. Suppose that the promise sequence and sequence of state vectors converge to steady-state values ω_{ss} and x_{ss} respectively, and that the multipliers on constraints (16) to (19) converge to $\lambda_{ss}^{h,0}(\sigma)$, $\lambda_{ss}^{h}(\sigma)$, $\lambda_{ss}^{g,0}(\sigma)$ and $\lambda_{ss}^{g,1}(\sigma)$ respectively. Then:*

1. *The promise sequence satisfies RWPE only if for all $\sigma' \in \Sigma$:*

$$\begin{bmatrix} \lambda_{ss}^{h}(\sigma') \\ \lambda_{ss}^{g,1}(\sigma') \end{bmatrix} = \beta \sum_{\sigma \in \Sigma} \frac{P(\sigma'|\sigma)\, P(\sigma)}{P(\sigma')} \begin{bmatrix} \lambda_{ss}^{h,0}(\sigma) + \lambda_{ss}^{h}(\sigma) \\ \lambda_{ss}^{g,0}(\sigma) \end{bmatrix}. \tag{23}$$

2. *If the promise-value function $V(\cdot, x)$ is quasi-concave for all $x \in X$, then the promise sequence satisfies RWPE if (23) holds.*

The results in Proposition 8 are analogous to a standard first-order condition. They provide a necessary condition for RWPE to hold, which becomes sufficient when the relevant objectives are quasi-concave. The measure $\frac{P(\sigma'|\sigma)P(\sigma)}{P(\sigma')}$ is a 'reverse' transition probability that the predecessor to an observed state σ' was state σ. In the social insurance example, σ is the number of time periods that have elapsed since the agent last received a high income draw. In this case, for any $\sigma' > 0$ it must be that $\sigma = \sigma' - 1$, so the probability that σ' was preceded by $\sigma' - 1$ is one and the probability that it was preceded by any other state is zero. When $\sigma = 0$ the agent has just received a high income draw, the contemporary participation constraint generally binds, and lagged multipliers are no longer relevant. The linear-quadratic inflation bias and capital tax examples are both fully deterministic, so for these the probabilistic term can be dropped. In all three examples, condition (23) implies a downward drift in the multiplier on promise-keeping constraints – the left side of the equation – relative to the multiplier on the relevant promise-making constraints and/or past promise keeping constraints – the right side. This drift occurs at the rate of pure time preference.

The downward drift in the multipliers contrasts with the outcome under Ramsey policy characterised by (21). This expression implies a steady-state multiplier recursion similar to (23), but without any downward drift. Instead, the coefficient on past promise-making constraints under Ramsey-optimal policy is fixed to one. This results in non-stationary evolution of the multipliers in the infinite-horizon case, which in turn can impart uncomfortable long-run properties such as the absence of long-run redistribution in the social insurance example. In the case of one-period ahead constraints, it implies equality between the shadow benefit of making a promise yesterday and the shadow cost of keeping it today. This likewise can imply surprising convergence results to Ramsey-optimal policy, such as the zero long-run inflation rate in the linear-quadratic inflation bias example and the zero long-run capital tax rate in the capital tax example. The next section shows that neither of these outcomes satisfies RWPE.

Proposition 8 relates solely to steady-state outcomes. It does not specify any properties of the transition dynamics associated with RWPE promise sequences. Indeed, if the value function is quasi-concave then it follows from the second part of Proposition 8 that if one promise sequence satisfies RWPE then any promise sequence that converges to the same ω_{ss} and induces convergence of the state vector to the same x_{ss} must also satisfy RWPE. This property holds regardless of quasi-concavity, though quasi-concavity is useful in confirming RWPE. The intuition is that weak Pareto efficiency becomes harder to satisfy over time. Initially, the only requirement is that there is no strict improvement for all policymakers from period 0 onwards. Eventually, however, convergence to steady state occurs and continued weak Pareto efficiency requires only that there should not exist strict gains relative to that steady state alone. If this is satisfied, then trivially it must be the case that no strict gain is available in earlier periods either.

6. Uniqueness and policy dynamics

Proposition 8 in the previous section proposes weak Pareto efficiency as a generalisation of optimality that can be applied recursively in Kydland and Prescott problems. The Proposition also suggests that transition dynamics will not be tied down by RWPE alone. In this section, two additional refinements are introduced that lead to uniqueness in the transition dynamics. The implications for policy are explored in each of our three examples. Formally, the need is to find a choice rule $C(x_0)$ that selects a unique dynamic promise sequence $\{\omega_s\}_{s=0}^{\infty}$ to be associated with any given initial state vector x_0, where $C : X \to \Omega(x_0)$. The choice rule will be assumed to nest RWPE, selecting only promise sequences that converge to a steady state where (23) holds. The question is what additional structure must be imposed on $C(x_0)$ to guarantee that transition dynamics are unique.

6.1. Time-invariance

The search for alternative approaches to policy design derives in part from the implausible dynamics that characterise Ramsey policy in a large range of different settings. In light of this, it is attractive to focus on choice rules that are time invariant in the following sense:

Assumption 6. *Suppose the promise sequence $\{\omega_s'\}_{s=0}^{\infty} = C(x_0)$ induces a sequence of state variables $\{x_{s+1}'\}_{s=0}^{\infty}$. **Time invariance** requires $C(x_t') = \{\omega_s'\}_{s=t}^{\infty}$ for all $t > 0$.*

In words, the continuation sequence of promises from period t onwards must be the same as that which would have been chosen were period t to be the first period and x_t' the initial state vector. For models without endogenous state variables, this assumption is equivalent to imposing that the promise vector should be constant in all periods. Combined with the requirement that RWPE should be satisfied, it is generally sufficient to deliver a unique path for policy variables in such models.[38] The implied allocation has a very desirable feature:

Proposition 9. *Suppose there are no state variables in the model. Then a choice rule satisfying RWPE and time invariance will induce an allocation $\{a_s\}_{s=0}^{\infty}$ such that $a_s = \bar{a}$ for all s and some $\bar{a} \in A$. This $\bar{a}$ is optimal in the set of feasible constant policies.*

The proof is immediate. Suppose a constant allocation were feasible and superior to a given choice rule satisfying RWPE. Associated with this constant allocation must be a constant sequence of promises that is time-invariant by construction.[39] But in each period a switch to the superior constant promises would be strictly preferred, a violation of RWPE.□

The set of models without state variables to which Proposition 9 relates is a relatively restrictive one, but it is a comforting endorsement of the RWPE criterion that it selects an allocation that can unambiguously be described as the best time-invariant choice in such settings. This is particularly

[38] In principle, there could be multiple time-invariant promise vectors satisfying the steady-state requirement for RWPE. But in models without endogenous state variables, steady state can be obtained immediately. It follows then from its definition that RWPE can only be satisfied by a subset of steady-state promise vectors that deliver the same value, otherwise switching from one constant vector to another would deliver a Pareto improvement. Hence there can be no multiplicity in the *value* of RWPE promises.

[39] These can be obtained by solving conditions (16) to (19) for the ω_s terms, given the allocation.

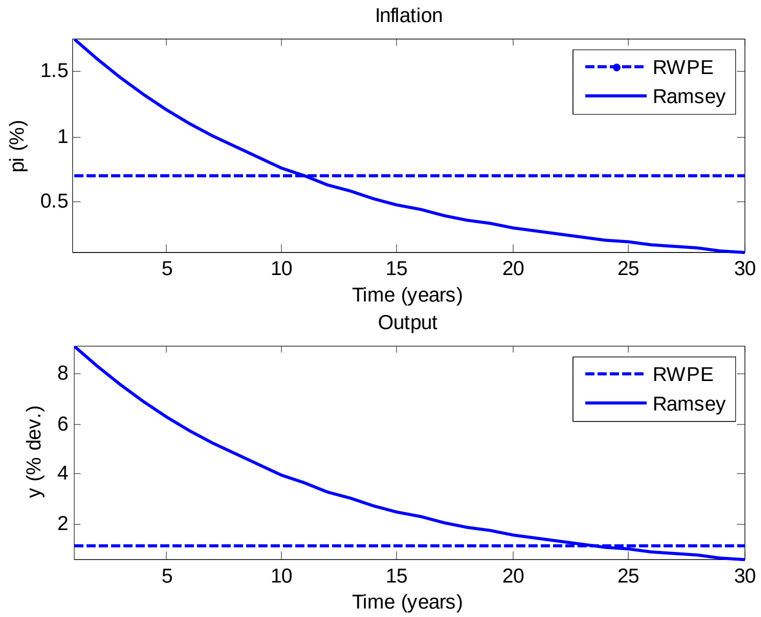

Figure 4: Ramsey and time-invariant RWPE policy in the inflation bias example

important, given that Woodford's (1999, 2003) 'timeless perspective' approach does not deliver the best time-invariant choice when it recommends that the policymaker should implement the Ramsey steady state policy immediately.

The implications of the time-invariant RWPE choice rule are next demonstrated in our two examples without endogenous state variables, the linear-quadratic inflation bias problem and the social insurance problem.

Example 1: A linear-quadratic inflation bias problem

Figure 4 plots the paths of inflation and output when the policy choice rule satisfies both RWPE and time-invariance in the simple linear-quadratic inflation bias example. Ramsey paths for inflation and output are superimposed for ease of comparison. There is a noticeable sense in which the time-invariant RWPE policy delivers a convex combination of the short-run and long-run outcomes of Ramsey policy. Time-invariant RWPE policy has inflation permanently above zero, which permits a small positive output gap to exist in perpetuity.[40]

The value of the welfare criterion under Ramsey and time-invariant RWPE policy is plotted in Figure 5. Neither policy Pareto-dominates the other in period 0, because early policymakers prefer the Ramsey policy whereas later policymakers prefer the RWPE policy.[41] But after 18 periods, the RWPE policy comes to strictly Pareto-dominate the continuation of the Ramsey policy. This reflects the fact that Ramsey policy does not satisfy recursive efficiency. If the timeless perspective approach

[40] Recall that the New Keynesian Phillips Curve (7) is not vertical in the long run, so permanently positive inflation is consistent with output being permanently above its flexible price level.

[41] Given the absence of endogenous state variables, the definition of Pareto efficiency here does not differ between *ex-ante* and *ex-post* approaches.

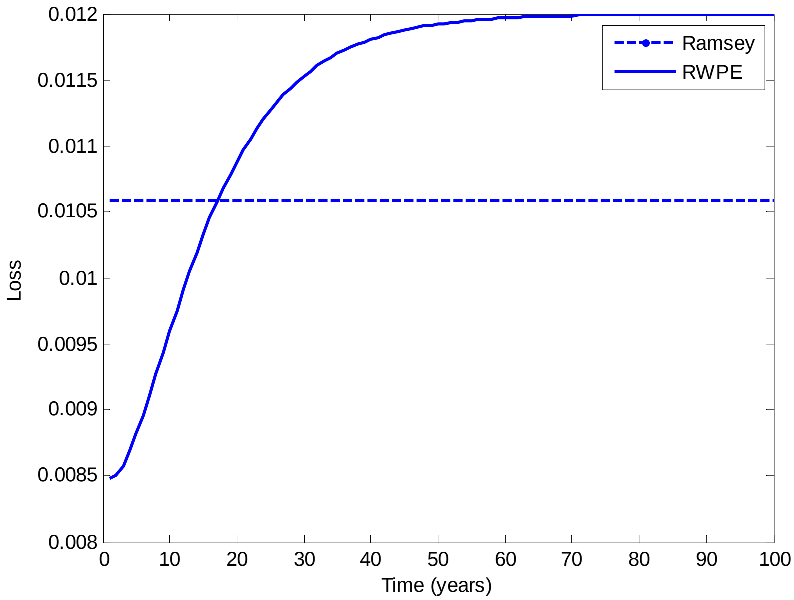

Figure 5: Losses under Ramsey and time-invariant RWPE policy in the inflation bias example

to policy design were to be adopted then the constant policy associated with Ramsey steady state would be implemented immediately, which clearly delivers a higher loss than our RWPE alternative.

Example 3: Social insurance with participation constraints

The Ramsey and time-invariant RWPE policies for the consumption of permanently low-income agents are compared in Figure 6. The time-invariant RWPE policy has agents with permanently low income consuming at a fixed constant level that is a little over 5 per cent above their endowment. This contrast sharply with the Ramsey policy, which requires low-income agents to consume no more than their endowment after a finite number of periods. The distribution of consumption across other agents under the RWPE policy is time-invariant, and the Pareto weights of agents who have received high income draws in the past follows a stationary process consistent with condition (23). The contrast with Ramsey policy is again sharp, since in this setting Ramsey policy implies that the Pareto weights of agents who have received high income draws in the past are non-stationary, consistent with condition (21).

More formally, the consumption of an agent who last received a high-income draw σ periods ago under the RWPE policy solves:

$$\left[u'\left(c\left(\sigma\right)\right)\right]^{-1} = \left(\frac{1+\lambda\left(\sigma\right)}{\eta}\right), \tag{24}$$

and the promise-keeping multiplier $\lambda\left(\sigma\right)$ satisfies:

$$\lambda\left(\sigma+1\right) = \beta\lambda\left(\sigma\right), \tag{25}$$

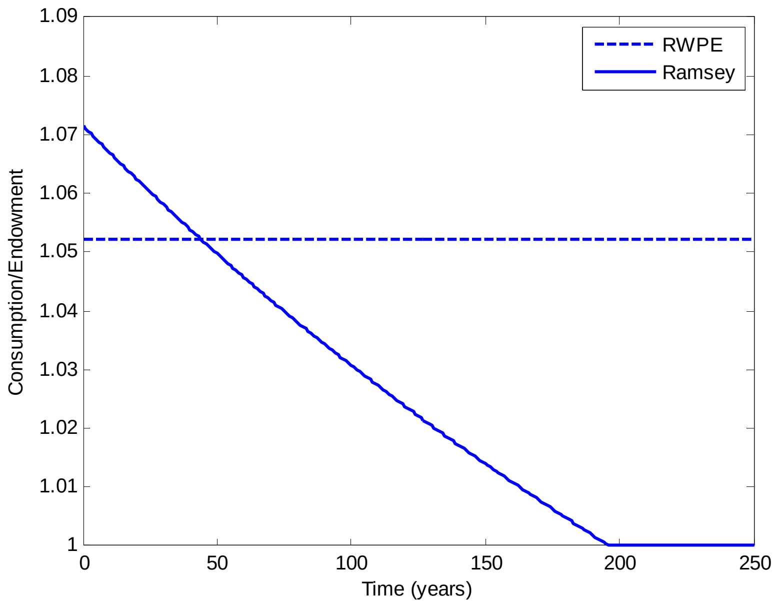

Figure 6: Ramsey and time-variant RWPE paths for consumption paths of low-income agents

for all $\sigma \geq 0$, where η is the time-invariant multiplier on the resource constraint. The component of the Pareto weight reflecting the policymaker's underlying utilitarian preferences – the 1 on the right side of (24) – has a permanent effect on consumption outcomes over time. By contrast, the Ramsey policy sees generic agent i receiving a Pareto weight of $1 + \lambda_t^i$ in period t, with the ratio $\lambda_t^i/(1 + \lambda_t^i)$ approaching 1 for almost all of the stochastic-income agents as time progresses.[42] This means that underlying social preferences have only a negligible impact on redistribution as t grows. Allocations in the limit become entirely dominated by the need to make good on past promises.

The welfare comparison between the two policies is given in Figure 7, based on the policymaker's utilitarian objective criterion. This figure charts the permanent, uniform level of consumption for all agents that would be equivalent in present-value welfare terms to continuation with each policy. It is normalised relative to the first-best, in which all agents consume the average endowment each period. Because the fraction of agents with permanently low incomes is relatively small, the absolute magnitude of losses relative to the first best is similarly small under both policies.[43] However, the Ramsey continuation plan is dominated by the time-invariant RWPE scheme after a relatively short time horizon of only 19 years.

6.2. Recursivity of the inner problem

Time invariance and RWPE in themselves are insufficient to deliver policy uniqueness when endogenous state variables are present. Whilst the steady state is tied down, an indeterminacy remains in

[42] This is true even for those agents whose income levels are permanently low, as eventually their participation constraints become binding and their consumption levels are driven down to their endowments.

[43] This is a model in which alternative symmetric welfare criteria, such as Rawlsianism, would deliver identical optimal policies. However, a Rawlsian would only place weight on the welfare of the low-income agents, so there would be greater variation in the objective criterion associated with each policy.

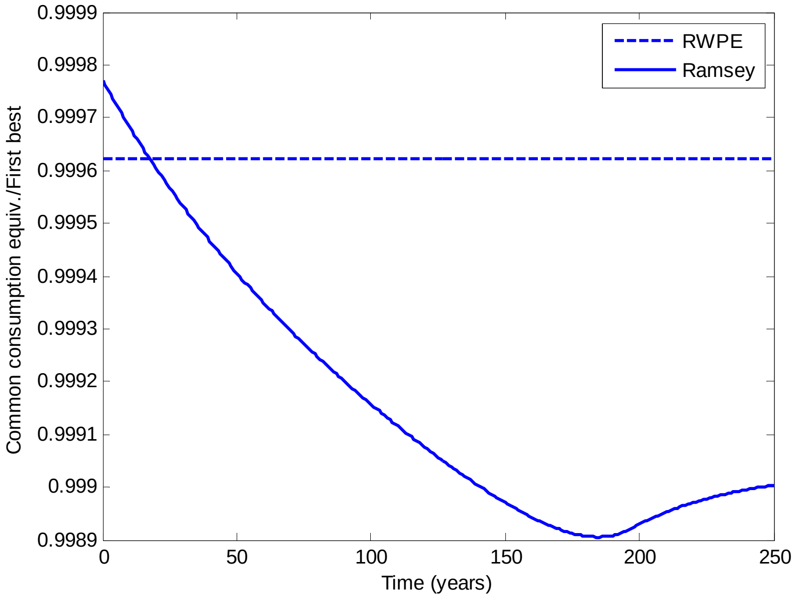

Figure 7: Welfare under Ramsey and RWPE policies in the social insurance example

the relationship between the endogenous state variables and promises, such that multiple transition dynamics are possible.

To make progress, it is useful to consider the promise-value function associated with a given initial state vector x_0, which can be denoted $V\left(C\left(x_0\right), x_0\right)$. If C induces a sequence of state vectors $\{x'_s\}_{s=0}^{\infty}$ then it follows from the time-consistency of the inner problem that for all $t \geq 0$:

$$V\left(C\left(x'_t\right), x'_t\right) = \max_{a_t, x_{t+1}} \left\{r\left(x'_t, a_t\right) + \beta V\left(C\left(x'_{t+1}\right), x_{t+1}\right)\right\}, \tag{26}$$

subject to constraints (2), (3) and (16) to (19) holding in period t. This is a standard recursive representation of the time-consistent component of the problem, given a choice for the promise sequence. If C satisfies time invariance then it follows that setting $x_{t+1} = x'_{t+1}$ must solve this problem, together with an appropriate choice of a_t.

An important feature of (26) is that the promise sequence $C(x'_{t+1})$ featuring in V on the right side is independent of the choice of x_{t+1}. Nonetheless, x'_{t+1} will by construction be an optimal choice of x_{t+1}, provided that C satisfies time invariance. In general, there is no reason why $V(C(x'_t), x'_t)$ needs to coincide with an alternative value function $\tilde{V}(C(x'_t), x'_t)$ that treats the promise sequence featuring in $\tilde{V}$ on the right hand side as *dependent* of the choice of x_{t+1}. Such an alternative value function is defined as a fixed point of the functional equation:

$$\tilde{V}\left(C\left(x_t\right), x_t\right) = \max_{a_t, x_{t+1}} \left\{r\left(x_t, a_t\right) + \beta \tilde{V}\left(C\left(x_{t+1}\right), x_{t+1}\right)\right\}, \tag{27}$$

for all $x_t \in X$, a given choice rule C, and subject to (2), (3) and (16) to (19) holding in period t.

Unlike V, the alternative value function $\tilde{V}$ allows the choice of x_{t+1} to impact on the set of future promise values through the choice rule $C\left(x_{t+1}\right)$. This contrasts with (26), where the future promise values are exogenous. When the set of future promises is treated as endogenous, a marginal change in x_{t+1} could induce a promise sequence that is strictly preferred from period $t+1$ onwards. Heuristically, it may be that $\frac{\partial \tilde{V}}{\partial C}\frac{\partial C}{\partial x_{t+1}} \neq 0$. This would change the choice of x_{t+1} in (27) relative to that in (26). The difference is really one of timing. Are promises fixed once-and-for-all at the beginning of time, or should they be updated in response to the endogenous state vector? Our final restriction asserts that the timing choice should make no difference.

Assumption. ***Recursivity of the inner problem*** *requires $C(x_t)$ is such that $V(C(x_t), x_t) = \tilde{V}(C(x_t), x_t)$ for all $x_t \in X$.*

A necessary condition for the inner problem to be fully recursive is that $\tilde{V}(C(x), x'_{t+1})$ must be at a local maximum when $x = x'_{t+1}$, taking the x'_{t+1} in the second argument in $\tilde{V}$ as fixed. If this is not the case, then there are marginal gains to changing the state variable away from x'_{t+1} under the alternative value function. These are captured in $\tilde{V}$ but not in V, which means the solutions to (26) and (27) would no longer coincide. This leaves the problem of finding a time-invariant and RWPE choice rule such that $\tilde{V}(C(x), x'_{t+1})$ will always be at a local maximum when $x = x'_{t+1}$. The simplest way of achieving this is to make $C(x)$ invariant in x and have a sequence of constant promises. The implication is that promises in every period are set equal to the steady-state values that are consistent with RWPE, and this is how we proceed below when analysing the capital tax example. We are open to the development of more nuanced approaches to the design of transition policy, but constant promises have the advantage of computational ease and relative transparency. In addition, it can be shown that *only* constant promises satisfy RWPE, time-invariance and recursivity of the inner problem in linear-quadratic environments.[44] Thus general alternatives do not appear available.

Example 2: A capital tax problem

The Ramsey and RWPE paths for capital taxes and the capital stock are illustrated in Figure 8.[45] In both cases, the capital stock is initially assumed to be equal to its Ramsey steady-state level. The evolution of capital taxes under the RWPE policy exhibits much more stability than under Ramsey, with rates gradually falling from an initial 42 per cent to a steady-state value of around 28 per cent. The fall in taxes tracks the decline in the capital stock, rather than being driven by dynamics intrinsic to the policy itself. This is unlike the Ramsey case, in which capital first falls in response to high taxes but then ultimately returns to its initial level.

The capital tax rate is always positive under RWPE, in an environment where capital taxes nonetheless converge to zero under Ramsey policy. We take this as challenging the notion that capital taxation is a 'bad idea' in the Judd model. There is no justification for immediately implementing zero capital taxes in this model. Our approach to policy design delivers taxes in and out of steady-state that are modestly positive, based on a recursive selection criterion. The Ramsey plan has very high transitional capital taxes, and these cannot be meaningfully separated from the long-run converge to zero.

[44] More precisely, a constant promise sequence is the only linear function of the endogenous state variables that satisfies RWPE, time-invariance and recursivity of the inner problem in environments that are linear-quadratic. A proof is available on request.

[45] The capital stock is charted as a proportion of the Ramsey steady-state level.

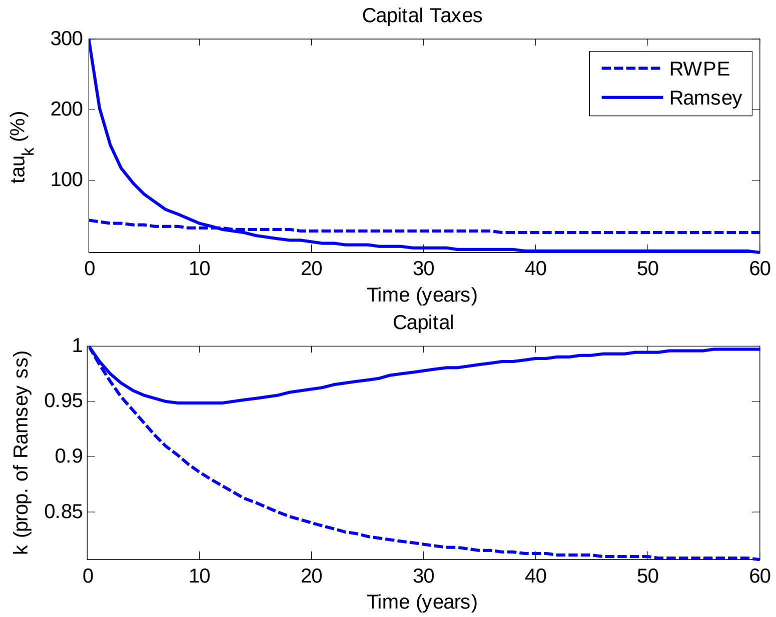

Figure 8: Capital taxes and the capital stock: Ramsey versus RWPE policy

To illustrate the dependence of RWPE capital tax rates on the capital stock alone, Figure 9 charts tax rates and the capital stock under differing assumptions about the initial value of k_0. It shows RWPE policy when capital starts 5 per cent above steady state, 5 per cent below steady state, and in steady state. When capital is below steady state, so are capital taxes. Likewise, when capital is above steady state, so are taxes. Under Ramsey policy, initial capital tax rates are implausibly high irrespective of k_0.

The RWPE criterion delivers zero steady-state taxes in the limiting case when the discount factor approaches unity. This is unsurprising. With a sufficiently high degree of patience, no policymaker would want to impede the accumulation of capital. What is distinctive in our approach is the ability to separate the appropriate treatment of the time-inconsistency problem from the question of long-run time preferences. Zero capital taxes are not a steady-state outcome of the Ramsey plan because the Ramsey policymaker has a preference for long-run outcomes. Instead, they are a by-product of optimally targeting welfare gains at the first generation when commitments in period 0 can be designed with complete freedom.

The separate roles of time inconsistency and long-run preferences become clearer when comparing the welfare levels associated with each policy. Welfare comparisons are more complex than before, as the long-run capital stock is important for welfare and the evolution of the capital stock differs between the Ramsey and RWPE policies. However, our focus is on an *ex-post* objective criterion based on the value of switching away from a chosen commitment. Figure 10 plots the value in a given period of switching from the RWPE policy to the continuation promise sequence associated with the Ramsey policy. That is, the figure shows the value of an unanticipated switch from the RWPE promises$\{\omega_s^P\}_{s=t}^{\infty}$ to the Ramsey continuation promises $\{\omega_s^R\}_{s=t}^{\infty}$ in period t, given a capital stock that has been induced by the RWPE promise sequence $\{\omega_s^P\}_{s=0}^{t}$ up to period t. The RWPE

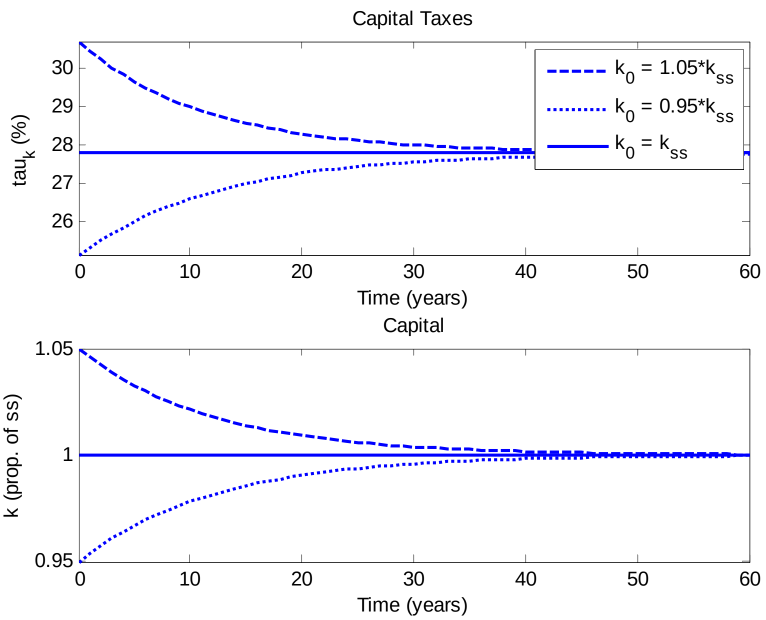

Figure 9: Capital taxes and the capital stock: the effect of initial conditions

policy is such that switching to any fixed alternative plan cannot be welfare-improving once steady state has been reached.

The welfare gains from switching to the Ramsey promise sequence have been exhausted by the 14th period of the model.[46] There would of course be a gain in every period from switching to a Ramsey plan that starts in that particular period, but clearly this is not something that could be done in every period. The purpose of the *ex-post* Pareto criterion is to ask whether there is a fixed alternative dynamic plan that is uniformly preferred to continuing with a pre-existing commitment. By definition, no such alternative plan exists relative to any policy that is weakly Pareto efficient. The only policy that satisfies this normative requirement recursively is RWPE.

7. Conclusion

Macroeconomists are frequently asked the basic normative question 'How *should* policy be designed?' The obvious response is that it should be chosen to be best according to an accepted social welfare criterion. The difficulty in Kydland and Prescott problems is that even given such a criterion, what is 'best' depends on the time period in which choice is viewed. In order to provide policy advice, we cannot avoid asking 'Best for whom?'

When commitment devices are present, as we have assumed, policy is allowed to be designed once-and-for-all in period zero. This means it is clearly possible to choose policy so that it is best for period zero, as Ramsey policy does. Since the economist is herself positioned in period zero when analysing the problem, it is tempting to view this as a correct response to 'Best for whom?'

Our paper has implicitly taken a different perspective. To justify this, it is useful to draw a parallel

[46]A similar point could be made by considering the gains from switching away from the Ramsey policy.

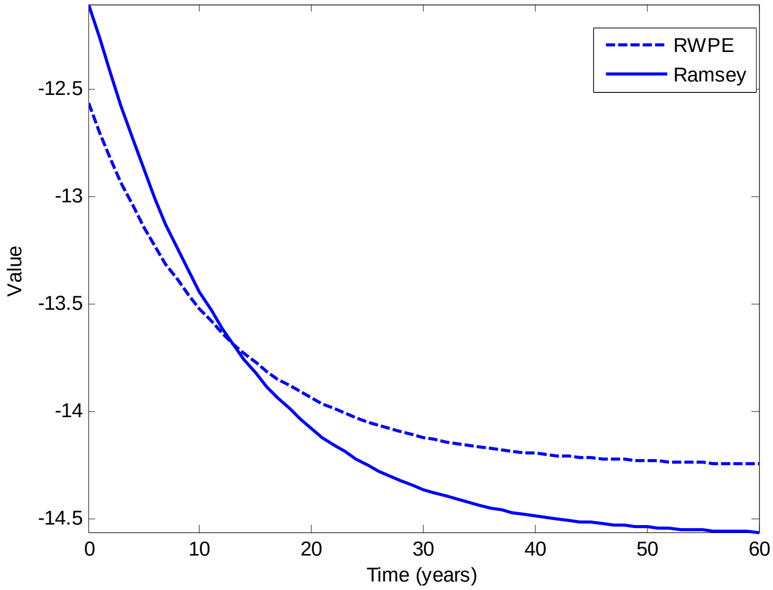

Figure 10: Welfare from switching from RWPE to Ramsey continuation promises in the capital tax example

with the classic social choice literature initiated by Arrow (1951). The focus of this literature was on devising general choice procedures for environments with multiple competing preferences. It was treated as axiomatic that dictatorship was an undesirable feature of a choice rule. This did not depend on whether Arrow himself shared the preferences of the proposed dictator. The point is that there may be more fundamental principles relating to appropriate social choice that are not reflected in the basic social welfare criterion. Just because the initial generation *can* impose its preferences on all subsequent generations does not mean that it *should*.

For this reason it is important to study alternative normative choice procedures to Ramsey policy in Kydland and Prescott problems. This paper provides one such alternative. As we have discussed at length, by retaining recursive applicability our Pareto criterion overcomes many of the implausible features associated with Ramsey policy. There is no policy transition independently of the economy's natural state vector, and long-run outcomes must be meaningfully desirable.

Though our focus has been normative, the policies that result have interesting parallels with recent positive analysis of commitment problems by Sleet and Yeltekin (2006) and Golosov and Iovino (2014). These authors focus on the best reputational equilibria that can be supported in dynamic models of asymmetric information, with no aggregate endogenous states. They show that the resulting policies are equivalent to Ramsey-optimal strategies for a policymaker whose discount factor exceeds the private sector's. This occurs because future generations' welfare must be given sufficient weight in order for the equilibrium to be sustained over time. Our recursive Pareto criterion similarly ensures that later generations do not inherit an excessive burden of past promises, though by a very different route. In models without state variables, it is consistent with policy that maximises steady-state welfare. In further work we hope to explore the relationship between policy that is recursively

Pareto efficient and policy that is sustainable in a reputational equilibrium.

Finally, the attention throughout this paper has been on settings with no aggregate risk. There is no intrinsic barrier to relaxing this, but doing so would add one extra degree of complexity to the problem. Policy preferences would not just differ between policymakers in different time periods, but also within each period, between policymakers in different states of the world. Some way would be needed to resolve this additional form of disagreement. One approach would be to impose a Rawlsian 'veil of ignorance' when institutional design takes place, so that current and past realisations of the shock process are known only probabilistically. This would allow the value of policies under different histories to be aggregated into a single, time-invariant objective. Mathematically this is what we are already doing in our social insurance example, where stochastic individual utilities are aggregated into a utilitarian social welfare function. The difference seems one of interpretation. We plan to explore this in further work.

References

[1] Ábrahám, Á., and S. Laczó (2014), 'Efficient Risk Sharing with Limited Commitment and Storage', manuscript.

[2] Abreu, D., D. Pearce and E. Stacchetti (1990), 'Toward a Theory of Discounted Repeated Games with Imperfect Monitoring', *Econometrica*, 58 (5), 1041–1063.

[3] Adam, K., and M. Woodford (2012), 'Robustly Optimal Monetary Policy in a Microfounded New Keynesian Model', *Journal of Monetary Economics*, 59(5), 468–487.

[4] Atkeson, A. (1991), 'International Lending with Moral Hazard and Risk of Repudiation', *Econometrica*, 59(4), 1069–1089.

[5] Atkeson, A., V.V. Chari and P.J. Kehoe (1999), 'Taxing Capital Income: A Bad Idea', *Federal Reserve Bank of Minneapolis Quarterly Review,* 23, 3-18.

[6] Atkeson, A., and R.E. Lucas Jr. (1992), 'On Efficient Distribution with Private Information', *Review of Economic Studies*, 59(3), 427–453.

[7] Arrow, K. (1951), *Social Choice and Individual Values*, New Haven: Yale University Press.

[8] Barro, R.J., and D.B. Gordon (1983), 'Rules, Discretion and Reputation in a Model of Monetary Policy', *Journal of Monetary Economics*, 12, 101–121.

[9] Benigno, P., and M. Woodford (2012), 'Linear-quadratic Approximation of Optimal Policy Problems', *Journal of Economic Theory*, 147(1), 1–42.

[10] Blake, A.P. (2001), 'A "Timeless Perspective" on Optimality in Forward-Looking Rational Expectations Models', Working Paper 188, NIESR

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

[12] Broer, T. (2013), 'The Wrong Shape of Insurance? What Cross-Sectional Distributions Tell Us about Models of Consumption Smoothing', *American Economic Journal: Macroeconomics*, 5(4), 107–140.

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

[14] Chari, V.V., and P.J. Kehoe (1990), 'Sustainable Plans', *Journal of Political Economy*, 98, 783–802.

[15] Chari, V.V., and P.J. Kehoe (1999), 'Optimal Fiscal and Monetary Policy', in J.B. Taylor and M. Woodford (eds.) *Handbook of Macroeconomics*, Vol. 1C, Amsterdam: Elsevier.

[16] Corsetti, G., L. Dedola and S. Leduc (2010), 'Optimal Monetary Policy in Open Economies', in B.M. Friedman and M. Woodford (eds.), *Handbook of Monetary Economics*, Vol. 3, Amsterdam: Elsevier.

[17] Damjanovic, T., V. Damjanovic, and C. Nolan. (2008), 'Unconditionally Optimal Monetary Policy', *Journal of Monetary Economics*, 55, 491–500.

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

[19] Ellison, M., and N. Rankin (2007), '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.

[20] Farhi, E., and I. Werning (2007), 'Inequality and Social Discounting', *Journal of Political Economy*, 115(3), 365–402.

[21] Farhi, E., and I. Werning (2010), 'Progressive Estate Taxation', *The Quarterly Journal of Economics*, 125(2), 635–673.

[22] Golosov, M., L.E. Jones and M. Tertilt (2007), 'Efficiency with Endogenous Population Growth', *Econometrica*, 75(4), 1039–1071.

[23] Golosov, M., and L. Iovino (2014), 'Social Insurance, Information Revelation, And Lack Of Commitment', manuscript.

[24] Hansen, L.P., and T.J. Sargent (2008), *Robustness*, Princeton: Princeton University Press.

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

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

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

[28] Kocherlakota, N.R. (1996), 'Implications of Efficient Risk Sharing without Commitment', *Review of Economic Studies*, 63(4), 595–609.

[29] Kocherlakota, N.R. (2010), *The New Dynamic Public Finance*, Princeton: Princeton University Press.

[30] Krueger, D., and F. Perri (2006), 'Does Income Inequality Lead to Consumption Inequality? Evidence and Theory', *Review of Economic Studies*, 73(1), 163–193.

[31] Krueger, D., and H. Uhlig (2006), 'Competitive risk sharing contracts with one-sided commitment', *Journal of Monetary Economics*, 53(7), 1661–1691.

[32] Kydland, F.E., and E.C. Prescott (1977), 'Rules Rather than Discretion: The Inconsistency of Optimal Plans', *Journal of Political Economy*, 85(3), 473–492.

[33] Ljungqvist, L., and T.J. Sargent (2012), *Recursive Macroeconomic Theory*, 3rd Edition, Cambridge (MA): MIT Press.

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

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

[36] Niemann, S., P. Pichler and G. Sorger (2013), 'Public debt, discretionary policy, and inflation persistence', *Journal of Economic Dynamics and Control*, 37(6), 1097–1109.

[37] Ortigueira, S. (2006), 'Markov-perfect optimal taxation', *Review of Economic Dynamics*, 9(1), 153–178.

[38] Phelan, C. (2006), 'Opportunity and Social Mobility', *Review of Economic Studies*, 73(2), 487–505.

[39] Phelan, C., and E. Stachetti (2001), 'Sequential Equilibria in a Ramsey Tax Model', *Econometrica*, 69(6), 1491–1518.

[40] Ramsey F.P. (1928), 'A Mathematical Theory of Saving', *Economic Journal*, 38, 543–559

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

[42] Rogerson (1985), 'Repeated Moral Hazard', *Econometrica*, 53(1), 69–76.

[43] Sleet, C., and S. Yeltekin (2006), 'Credibility and Endogenous Societal Discounting', *Review of Economic Dynamics*, 9(3), 410–437.

[44] Stern, N. (2006), *The Economics of Climate Change: The Stern Review*, Cambridge: Cambridge University Press

[45] Straub, L., and I. Werning (2014), 'Positive Long Run Capital Taxation: Chamley-Judd Revisited', manuscript.

[46] Svensson, L.E.O. (1999), 'How Should Monetary Policy Be Conducted In An Era Of Price Stability?', CEPR Discussion Paper 2342.

[47] Thomas, J., and T. Worrall (1990), 'Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem', *Journal of Economic Theory*, 51, 367–9.

[48] Wachsmuth, G. (2013), 'On LICQ and the Uniqueness of Lagrange Multipliers', *Operations Research Letters*, 41(1), 78–80.

[49] 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.

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

[51] Woodford, M. (2010), 'Optimal Monetary Stabilization Policy', in B.M. Friedman and M. Woodford (eds.), *Handbook of Monetary Economics*, Vol. 3, Amsterdam: Elsevier.

A. Proofs

A.1. Proof of Proposition 1

Suppose that the statement were not true. Then for some $\tau \geq 1$ there must exist a continuation allocation $\left\{x''_{s+1}, a''_s\right\}_{s=t+\tau}^{\infty}$ that satisfies the constraints of the inner problem from τ onwards and delivers a higher value $W_{t+\tau}$ than the allocation $\left\{x'_{s+1}, a'_s\right\}_{s=t+\tau}^{\infty}$. But if this is true then the combined allocation $\left\{\left\{x'_{s+1}, a'_s\right\}_{s=t}^{t+\tau-1}, \left\{x''_{s+1}, a''_s\right\}_{s=t+\tau}^{\infty}\right\}$ must in turn deliver a higher value for W_t than $\left\{x'_{s+1}, a'_s\right\}_{s=t}^{\infty}$. By the optimality of $\left\{x'_{s+1}, a'_s\right\}_{s=t}^{\infty}$ this cannot be possible unless the combined allocation violates one of the constraints: (2), (3), or (16) to (19) in some period $s \geq t$. The component $\left\{x'_{s+1}, a'_s\right\}_{s=t}^{t+\tau-1}$ must satisfy these constraints from t to $t+\tau-1$, independently of the continuation outcome from $t+\tau$ onwards, since it is a part of the feasible sequence $\left\{x'_{s+1}, a'_s\right\}_{s=t}^{\infty}$, and for a given promise sequence the constraint set from t to $t+\tau-1$ is independent of outcomes from $t+\tau$ onwards. By assumption the continuation $\left\{x''_{s+1}, a''_s\right\}_{s=t+\tau}^{\infty}$ satisfies all such constraints for $t+\tau$ onwards, so we have a contradiction.

A.2. Proof of Proposition 2

The spaces $\Omega\left(x_t\right)$ and $\left(A \times X\right)^{\infty}$ contain bounded sequences, and thus are l^{∞} metric spaces when endowed with the sup norm. Quasi-concavity of the functions h_0, h, g_0 and g_1 together with the linear presence of the promise values in (16) to (19) implies that the constraint set for the inner problem is continuous in $\{\omega_s\}_{s=t}^{\infty}$ at all interior points of $\Omega\left(x_t\right)$. The objective function r is continuous under Assumption 1, so the result follows by a standard application of Berge's Theorem of the Maximum.

A.3. LICQ and Proof of Proposition 3

LICQ

Fix x_t and $\{\omega_s\}_{s=t}^{\infty}$. Let the allocation $\left\{x'_{s+1}, a'_s\right\}_{s=t}^{\infty}$ solve the associated inner problem. Fix a time period $s \geq t$, and stochastic draw σ_s, and suppose that in period s, q of the $(2j+2k)$ constraints in (16) to (19) are binding. Clearly $q \leq m$, where the right-hand side of this inequality is the dimensionality of $A(\sigma)$. Denote by H the $q \times m$ matrix of derivatives of the *binding* constraint functions in (16) to (19). That is, H takes the form:

$$H := \begin{bmatrix} \frac{\partial h_0^1}{\partial a_s^1} & \frac{\partial h_0^1}{\partial a_s^2} & \cdots & \frac{\partial h_0^1}{\partial a_s^m} \\ \frac{\partial h_0^2}{\partial a_s^1} & \frac{\partial h_0^2}{\partial a_s^2} & \cdots & \frac{\partial h_0^2}{\partial a_s^m} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial h^1}{\partial a_s^1} & \frac{\partial h^1}{\partial a_s^2} & \cdots & \frac{\partial h^1}{\partial a_s^m} \\ \vdots & \vdots & \ddots & \vdots \end{bmatrix}$$

where the superscripts on the h_0 and h functions index only the subset of constraints of types (16) and (17) that bind. The derivatives of the g_0 and g_1 functions would enter the H matrix likewise.

We say that LICQ is satisfied for x_t and $\{\omega_s\}_{s=t}^{\infty}$ when the matrix H is of full rank q for all $s \geq t$.

Proof of Proposition 3

Strict concavity in the return function r and the convexity of the constraint set implied by Assumptions 3 and 4 implies that the solution to the inner problem is unique at all points $\{\omega_s\}_{s=t}^{\infty} \in \mathring{\Omega}(x_t)$. Applying Berge's Theory of the Maximum, this solution must be continuous in $\{\omega_s\}_{s=t}^{\infty}$.[47] Associated with this solution is a standard set of Kuhn-Tucker conditions. Within period s and for stochastic draw σ_s, these conditions with respect to a_s will take the form:

$$\frac{\partial r}{\partial a_s} + \lambda'_s(\sigma_s) \frac{\partial G(a_s(\sigma_s))}{\partial a_s} = 0 \tag{28}$$

$$\lambda'_s(\sigma_s) G(a_s(\sigma_s)) = 0 \tag{29}$$

$$\lambda'_s(\sigma_s) \geq 0 \tag{30}$$

where $G(a_s(\sigma_s))$ stacks the constraints in (16) to (19):

$$G(a_s(\sigma_s)) := \begin{bmatrix} h_0(a_s(\sigma_s)) + \beta \mathbb{E}_s \omega_{s+1}^h(\sigma_{s+1}) \\ h(a_s(\sigma_s)) + \beta \mathbb{E}_s \omega_{s+1}^h(\sigma_{s+1}) - \omega_s^h(\sigma_s) \\ g_0(a_s(\sigma_s)) + \beta \mathbb{E}_s \omega_{s+1}^g(\sigma_{s+1}) \\ g_0(a_s(\sigma_s)) - \omega_s^g(\sigma_s) \end{bmatrix}$$

[47] C.f. the proof of Proposition 2.

and $\lambda_s = \left[\left(\lambda_s^{h,0}\right)', \left(\lambda_s^{h}\right)', \left(\lambda_s^{g,0}\right)', \left(\lambda_s^{g,1}\right)'\right]'$ is the stacked vector of multipliers on these constraints.[48] The only barrier to asserting an envelope condition is to guarantee the existence and uniqueness of these multipliers. This is a straightforward matrix invertability problem, and is indeed ensured so long as LICQ holds: see, for instance, Wachsmuth (2013). A standard limiting argument can then establish that the envelope theorem applies in this case, and the derivatives of the value function with respect to ω_s are as given in the Proposition.

A.4. Proof of Proposition 4

For simplicity we drop explicit dependence on σ_s in this proof. Consider two promise sequences $\{\omega_s'\}_{s=t}^{\infty}, \{\omega_s''\}_{s=t}^{\infty} \in \Omega(x_t)$ such that $V\left(\{\omega_s'\}_{s=t}^{\infty}, x_t\right) = V\left(\{\omega_s''\}_{s=t}^{\infty}, x_t\right) := \bar{V}$. To establish quasi-concavity we must show:

$$V\left(\left\{\alpha\omega_s' + (1-\alpha)\,\omega_s''\right\}_{s=t}^{\infty}, x_t\right) \geq \bar{V} \tag{31}$$

for all $\alpha \in (0,1)$.

Let $\mathbf{y}' := \left\{x_{s+1}', a_s'\right\}_{s=t}^{\infty}$ and $\mathbf{y}'' := \left\{x_{s+1}'', a_s''\right\}_{s=t}^{\infty}$ solve the inner problems associated with $\{\omega_s'\}_{s=t}^{\infty}$ and $\{\omega_s''\}_{s=t}^{\infty}$ respectively. It follows from the concavity of r (Assumption 4) that (31) must be satisfied provided the convex combination $\alpha\mathbf{y}' + (1-\alpha)\,\mathbf{y}''$ is feasible when the promise sequence is $\{\alpha\omega_s' + (1-\alpha)\,\omega_s''\}_{s=t}^{\infty}$. In this case $\alpha\mathbf{y}' + (1-\alpha)\,\mathbf{y}''$ will deliver at least $\bar{V}$, which is then a lower bound on $V\left(\{\alpha\omega_s' + (1-\alpha)\,\omega_s''\}_{s=t}^{\infty}, x_t\right)$. In the event that r is strictly concave, (31) would be satisfied strictly, and the value function will be strictly quasi-concave.

The linearity of l and the concavity of p respectively imply that if (2) and (3) are satisfied in all time periods by both $\mathbf{y}'$ and $\mathbf{y}''$ then they must also be satisfied by $\alpha\mathbf{y}' + (1-\alpha)\,\mathbf{y}''$. These constraints are unaffected by variations in the promise values. It remains only to show that constraints (16) to (19) are also satisfied. Consider (16). We need:

$$h_0\left(\alpha a_s' + (1-\alpha)\,a_s''\right) + \beta\left[\alpha\omega_{s+1}^h + (1-\alpha)\,\omega_{s+1}^h\right] \geq 0$$

Since the constraint is satisfied by both $\mathbf{y}'$ and $\mathbf{y}''$, we have:

$$\alpha h_0\left(a_s'\right) + (1-\alpha)\,h_0\left(a_s''\right) + \beta\left[\alpha\omega_{s+1}^h + (1-\alpha)\,\omega_{s+1}^h\right] \geq 0$$

But by concavity:

$$h_0\left(\alpha a_s' + (1-\alpha)\,a_s''\right) \geq \alpha h_0\left(a_s'\right) + (1-\alpha)\,h_0\left(a_s''\right)$$

establishing the desired inequality. An identical argument confirms that (17) to (19) are likewise satisfied for all $s \geq t$. This establishes the feasibility of $\alpha\mathbf{y}' + (1-\alpha)\,\mathbf{y}''$ when the promise sequence is $\{\alpha\omega_s' + (1-\alpha)\,\omega_s''\}_{s=t}^{\infty}$, completing the proof.

[48] To ease on notation we assume in writing (28) that there are no constraints of the form (2) or (3). These would not change the arguments: they would simply imply the addition of an extra set of terms independent of λ_s' in (28).

A.5. Proof of Proposition 6

Suppose instead that the Ramsey plan were a time-consistent optimal choice of promises, and the constraints associated with these promises were binding in some time period. We will treat the case in which the binding constraints are (18) and/or (19): symmetric logic could be applied to a case in which (16) and (17) bind. Let ω_τ^g be a promise vector that constrains the inner problem, and denote by $\lambda_s^{g,0}$ and $\lambda_s^{g,1}$ the vectors of current-value multipliers on constraints (18) and (19) respectively for generic period s.[49] These multipliers are strictly positive whenever the respective constraints bind. The marginal effect of increasing ω_τ^g on $V\left(\left\{\omega_s^R\right\}_{s=t}^{\infty}, x_t\right)$ is $\beta^{\tau-t}\left(\lambda_{\tau-1}^{g,0} - \lambda_\tau^{g,1}\right)$ for $\tau > t$, and $-\lambda_\tau^{g,1}$ for $\tau = t$, as shown in Proposition 3. Suppose first that the only binding constraint is (18), in period $\tau - 1$. Then the derivative of $V\left(\left\{\omega_s^R\right\}_{s=t}^{\infty}, x_t\right)$ with respect to some elements of the promise vector must be strictly positive for all $t < \tau$, contradicting optimality in these periods.[50] Suppose instead that constraint (19) binds in period τ. Then the derivative of $V\left(\left\{\omega_s^R\right\}_{s=t}^{\infty}, x_t\right)$ with respect to some elements of the promise vector must be strictly negative for $t = \tau$, contradicting optimality in this period.

A.6. Proof of Proposition 7

For simplicity we will assume that the value function is differentiable. From Proposition 6 we know that if the Ramsey-optimal policy is not time-consistent, some promise constraints of the form (16) to (19) must bind for all $\{\omega_s\}_{s=0}^{\infty} \in \Omega(x_0)$.[51] Fix some promise sequence $\{\omega_s'\}_{s=0}^{\infty} \in \Omega(x_0)$. We will show that $\{\omega_s'\}_{s=0}^{\infty}$ cannot satisfy recursive strict Pareto efficiency.

As in the proof of Proposition 6, we will assume that the binding constraints are all of the 'one-period' type (18) and/or (19). Symmetric logic can be applied when infinite-horizon constraints additionally bind. If a promise sequence is recursively weakly Pareto efficient then there cannot exist a bounded array of marginal changes to the promise vectors:[52]

$$\left\{\left\{\frac{d\omega_s(\sigma_s)}{d\theta}\right\}_{\sigma_s \in \Sigma}\right\}_{s=t}^{\infty}$$

such that the effect of these marginal changes on the promise-value function is positive for all $s \geq t$, strictly so for at least one s. Assessed in period s, from the envelope condition this effect is given by Δ_s:

$$\Delta_s := \sum_{\tau=s}^{\infty} \beta^{\tau-s} \sum_{\sigma_\tau \in \Sigma} P(\sigma_\tau) \left[\beta \lambda_\tau^{g,0}(\sigma_\tau) \sum_{\sigma_{\tau+1} \in \Sigma} P(\sigma_{\tau+1}|\sigma_\tau) \frac{d\omega_{\tau+1}(\sigma_{\tau+1})}{d\theta} - \lambda_\tau^{g,1}(\sigma_\tau) \frac{d\omega_\tau(\sigma_\tau)}{d\theta}\right]$$

where $P(\sigma_{\tau+1}|\sigma_\tau)$ is the transition probability between σ_τ and $\sigma_{\tau+1}$. Suppose first that only promise-

[49] We suppress dependence on the stochastic process σ to ease on notation.

[50] Note that increasing the promise vector expands the binding constraint set for the inner problem in this case, so is a movement within $\Omega(x_t)$.

[51] Suppose otherwise. The solution to the inner problem when promise constraints do bind must deliver lower value than the solution when they do not. The optimal policy without these constraints must therefore be a time-consistent Ramsey optimum.

[52] θ is an arbitrary parameter normalising the scale of the differential change.

making constraints ever bind, and that this is true in some period $\tau \geq t$, so that $\lambda_{\tau}^{g,0}(\sigma_\tau) > 0$ for some σ_τ. Since promise-keeping constraints do not bind, $\lambda_{\tau}^{g,1}(\sigma_\tau) = 0$ for all σ_τ. Then we can set $\Delta_s > 0$ for all $t \leq s \leq \tau$ by fixing $\frac{d\omega_{\tau+1}(\sigma_{\tau+1})}{d\theta} = 1$ across states in $\tau + 1$. This relaxes the promise-making constraint, and promise-keeping constraints do not bind at the margin, so it must represent a feasible differential movement within $\Omega(x_s)$ for all relevant s. $\Delta_s = 0$ for $s > \tau$, so the original promise sequence violates strict Pareto efficiency at t.

Now suppose instead that some promise-keeping constraints bind in state σ_t of period t. Then we can set $\Delta_t > 0$ by fixing $\frac{d\omega_t(\sigma_t)}{d\theta} = -1$, with all other promise changes fixed zero. This implies $\Delta_s = 0$ for $s > t$, and again we have a violation of strict Pareto efficiency at t.

A.7. Proof of Proposition 8

Part 1

As in the proof of Proposition 7, we will assume that all binding promise constraints are of the 'one-period' form (18) and (19) for simplicity: analogous reasoning works for the infinite-period constraint form.

A necessary condition for weak Pareto efficiency is that there should not exist a bounded sequence of marginal changes to the promise vectors:

$$\left\{ \left\{ \frac{d\omega_s(\sigma_s)}{d\theta} \right\}_{\sigma_s \in \Sigma} \right\}_{s=t}^{\infty}$$

such that the value of the object Δ_s, defined in the proof of Proposition 7, is either bounded above zero for all $s \geq t$ or bounded below zero for all $s \geq t$, for some $t \geq 0$. Consider the sequence that sets $\frac{d\omega_s(\sigma')}{d\theta} = 1$ for some $\sigma' \in \Sigma$ and all $s \geq t$, with $\frac{d\omega_s(\sigma)}{d\theta} = 0$ for $\sigma \neq \sigma'$ and all s. The value of Δ_s associated with this marginal change to the promise vector can be denoted $\Delta_s(\sigma')$:

$$\Delta_s(\sigma') = P(\sigma') \sum_{\tau=s}^{\infty} \beta^{\tau-s} \left[\sum_{\sigma \in \Sigma} \frac{P(\sigma'|\sigma) P(\sigma)}{P(\sigma')} \beta \lambda_{\tau}^{g,0}(\sigma) - \lambda_{\tau}^{g,1}(\sigma') \right]$$

For weak Pareto efficiency to hold at the limit as t becomes large, and thus for it to hold recursively, it is necessary for the limiting value of $\Delta_s(\sigma')$ to be zero. This implies:

$$\beta \sum_{\sigma \in \Sigma} \frac{P(\sigma'|\sigma) P(\sigma)}{P(\sigma')} \lambda_{ss}^{g,0}(\sigma) - \lambda_{ss}^{g,1}(\sigma') = 0$$

for all $\sigma' \in \Sigma$. Applying symmetric logic to cases in which infinite-horizon constraints bind, condition (23) follows.

Part 2

By usual logic, if the promise-value functions are quasi-concave in promises, then the absence of a local (differential) change to the promise sequence $\{\omega_s\}_{s=t}^{\infty}$ that is strictly preferred by all policymakers from t onwards implies the absence of a *global* strict Pareto improvement. Hence in this case it is sufficient to show that if (23) holds, there cannot exist a bounded sequence of marginal changes to the promise vectors:

$$\left\{ \left\{ \frac{d\omega_s(\sigma_s)}{d\theta} \right\}_{\sigma_s \in \Sigma} \right\}_{s=t}^{\infty}$$

such that the associated $\Delta_s \geq \varepsilon > 0$ for all $s \geq t$. Suppose otherwise. It follows from convergence of the multipliers and the boundedness of the marginal changes that for all s above some finite threshold and some $\sigma' \in \Sigma$ we must have:

$$\sum_{\tau=s}^{\infty} \beta^{\tau-s} \left[\beta \frac{d\omega_{\tau+1}(\sigma')}{d\theta} \sum_{\sigma_\tau \in \Sigma} \frac{P(\sigma'|\sigma) P(\sigma)}{P(\sigma')} \lambda_{ss}^{g,0}(\sigma) - \frac{d\omega_\tau(\sigma')}{d\theta} \lambda_{ss}^{g,1}(\sigma') \right] > \delta$$

for some $\delta > 0$. We know from (23) that:

$$\beta \sum_{\sigma \in \Sigma} \frac{P(\sigma'|\sigma) P(\sigma)}{P(\sigma')} \lambda_{ss}^{g,0}(\sigma) - \lambda_{ss}^{g,1}(\sigma') = 0$$

and the inequality in turn implies $\lambda_{ss}^{g,1}(\sigma') > 0$. Dividing through the inequality by $\lambda_{ss}^{g,1}(\sigma')$ and rearranging gives:

$$\sum_{\tau=s}^{\infty} \beta^{\tau-s} \frac{d\omega_{\tau+1}(\sigma')}{d\theta} > \sum_{\tau=s}^{\infty} \beta^{\tau-s} \frac{d\omega_\tau(\sigma')}{d\theta} + \frac{\delta}{\lambda_{ss}^{g,1}(\sigma')}$$

for all s sufficiently large. Now, define $\Gamma_s := \sum_{\tau=s}^{\infty} \beta^{\tau-s} \frac{d\omega_{\tau+1}(\sigma')}{d\theta}$. Boundedness of the sequence $\left\{ \frac{d\omega_s(\sigma')}{d\theta} \right\}_{s=t}^{\infty}$ implies the sequence $\{\Gamma_s\}_{s=t}^{\infty}$ likewise has a finite upper bound, say $\bar{\Gamma}$. This is inconsistent with the previous inequality: we have a contradiction.