Time-Consistent Institutional Design*

Charles Brendon[†] Martin Ellison[‡]

July 8, 2015

`Update in progress -- comments welcome`

This paper reconsiders normative policy design in environments subject to time inconsistency problems, à la Kydland and Prescott (1977). In these environments there are generally gains to making binding promises about future policy conduct, but decision-makers at different points in time have different preferences about the best promises to make and to keep. We consider the implications of using a Pareto criterion to resolve these differences. This criterion has the advantage that it can be satisfied recursively, in environments where, by definition, no policy is recursively *optimal*. It provides a basis for devising policies that are normatively appealing, but do not exhibit the sort of time variation that is a common property of Ramsey policy. We characterise outcomes in a general setting when promises satisfy the recursive Pareto criterion, and prove necessary conditions for generic existence of such policies. Three examples, explored throughout the paper, highlight how this approach can deliver simple, time-invariant rules for desirable policy conduct.

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, Vasco Carvalho, Giancarlo Corsetti, Tatiana Damjanovic, Paul Levine, Albert Marcet, Ramon Marimon, Alex Mennuni, Meg Meyer, Michel Normandin, Victor Ríos-Rull, 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, period-by-period discretionary policymaking will generically be sub-optimal. This is because successive generations of decisionmakers fail to endogenise the effects of their current choices on past expectations.[1]

A general feature of these settings is that all generations can be made better off if some way can be found to commit to binding policy promises.[2] A binding inflation target, for instance, can imply lower inflation expectations and a permanently improved inflation-output trade-off. A binding promise to limit future wealth taxation can ensure higher contemporary savings. To be effective these promises do not have to eliminate completely the scope for day-to-day policy choice, but they should restrict the options available. An inflation-targeting central bank retains operational independence, but is constrained by its obligation to target inflation.

The current paper takes as given that this sort of institutional commitment is possible, and revisits the policy question that it implies: How should such promises be designed? This is a non-trivial normative problem, because each generation of policymakers will have different preferences about the appropriate promises to make, and to keep, over time. Every generation would prefer that they themselves should be entirely free from promise-keeping obligations, but that their successors should be meaningfully restricted.

In this context it is impossible to find a sequence of promises that is recursively optimal: best for the first policymaker to choose, and best for all subsequent policymakers to adhere to. Choice must instead satisfy some weaker notion of desirability. The standard approach in the optimal policy literature is to relax recursivity, and to focus on the choice of promises that is optimal for the model's first time period. This is commonly known as Ramsey policy.

This paper sets out an alternative approach. We define and characterise policy when promises are not chosen to be 'best' in any one period, but must satisfy a meaningful desirability criterion at all times. In direct contrast with Ramsey policy, therefore, we retain recursivity in the choice procedure, but relax optimality. The criterion that we use in place of optimality is a version of the standard Pareto principle, taken with respect to the varying preferences of successive generations over promises. A choice satisfies the Pareto criterion if, when this choice is implemented, there is no other promise sequence that all current and future policymakers would prefer to switch to. It satisfies the criterion recursively if this statement remains true indefinitely, as the current period advances.

Our main motivation for developing this choice procedure is that a number of different literatures have found Ramsey policies to be unappealing. Two distinct issues arise. First, in models of monetary and fiscal policy it is common for Ramsey promises to induce very different policies in early time periods from later, even though the underlying state of the economy may not have changed at all.

[1] We refer to policymakers choosing in different time periods as belonging different 'generations', though in practice these time periods may be arbitrarily close.

[2] The formal insight that promises can be used to improve inefficient outcomes derives from the influential work of Abreu, Pearce and Stachetti (1990) on repeated games. There have been many applications in the macroeconomics literature, including Kocherlakota (1996), Chang (1998) and Phelan and Stachetti (2001).

Intuitively this is because Ramsey policy will never be constrained by inherited promises in the first time period, but will subsequently. The associated transition dynamic makes it hard to infer simple rules of best conduct that could be used to design mandates for policymaking institutions. The Ramsey-optimal inflation target, for instance, is likely to vary from one time period to another even when the economy's state vector does not. This problem has received particular attention in the New Keynesian literature, where Woodford (1999, 2003) has expressed the need for a 'timeless' approach to policy design, linking promises only to the state of the economy.

The second unappealing feature of Ramsey policy is that it can imply long-run outcomes that are extremely undesirable from the perspective of later generations. This is most commonly observed in the dynamic social insurance literature. In a dynamic hidden action model, Thomas and Worrall (1990) first showed that a Ramsey planner would promise to drive the consumption of almost all agents to zero over time. This 'immiseration' result was extended to general equilibrium in a hidden information setting by Atkeson and Lucas (1992), and has generated significant attention in the recent dynamic public finance literature.[3] Again, it obtains even though the state of the economy may be precisely the same in the long run as at the start of time. Influential contributions by Phelan (2006) and Farhi and Werning (2007) have explored alternative normative criteria that would prevent the result from going through, in both cases by reducing the extent to which future generations' welfare is discounted.

The recursive Pareto approach that we present avoids these two unappealing features of Ramsey policy by design. First, because the Pareto criterion must continue to apply in every time period, for a given state of the economy there is no reason why policy should take a different form in early periods from later. Thus the policy transition can be avoided. Second, because the choice of promises must satisfy the Pareto principle indefinitely, outcomes that are extremely undesirable for later generations will not arise.

Our paper defines and characterises the recursive Pareto principle as a way to choose promises. We find general conditions that guarantee the existence of a policy satisfying the recursive Pareto criterion, and illustrate the principle via three worked examples drawn from different branches of the optimal policy literature. These are, first, a linear-quadratic New Keynesian inflation bias problem; second, a variant of the Chamley (1986) and Judd (1985) capital tax problem; and, third, a dynamic social insurance problem under limited commitment, in the style of Kocherlakota (1996). To keep the arguments simple we focus on environments without aggregate risk.

An important feature of our work is the central focus that we place on promises as objects of analysis. This follows from a novel decomposition of generic Kydland and Prescott problems that we provide. We split the policy choice procedure into two components. The first, an 'inner problem', considers the best way to select the day-to-day aspects of policy, given an exogenous sequence of promises that must be kept each time period. We show that this inner problem is fully time consistent. It concerns the aspects of choice that remain once expectations have been fixed. The second component, the 'outer problem', considers how the promises themselves should be chosen. This corresponds to the practical problem of designing an institutional mandate. Under conventional restrictions the solution to the inner problem defines a well-behaved indirect utility function over the

[3] Kocherlakota (2010) provides a thorough discussion of immiseration in a dynamic Mirrlees setting.

set of possible current and future promises. That is, it defines preferences over promises. These preferences will differ from one period to the next, as the benefits from issuing promises are superseded by the costs of keeping them. We consider applying a recursive Pareto criterion to resolve this intertemporal difference in preferences.

The principal advantage of this decomposition is that it isolates the time-inconsistent aspects of choice for special treatment. It is only promises that are subject to our Pareto choice criterion, since it is only the choice of promises that is time-inconsistent. In environments that do not exhibit any time inconsistency there will be no role for promises, and our method will, trivially, deliver standard optimal policies. As we discuss below, this is not true of some alternative proposals in the literature.

Our most instructive general characterisation result, presented in Section 5, relates to the steady-state properties of policies that satisfy the recursive Pareto principle. We show that this steady state, if it exists, will differ systematically from the steady state of Ramsey policy. Specifically, less weight will be put on the value of manipulating past expectations. This is true in a precise sense. Whereas under Ramsey policy the shadow value associated with past promises will generally be a non-stationary object,[4] when policy satisfies a recursive Pareto criterion it decays over time at rate β – the discount factor.

Under slightly more restrictive assumptions we additionally prove a sufficiency result: any promises inducing convergence to a steady state of the form that we characterise must satisfy the recursive Pareto criterion. One implication of this is that the recursive Pareto criterion will usually be satisfied by multiple promise sequences. In Section 6 we explore ways to resolve this multiplicity in the context of our three main examples. In each case it is possible to find a time-invariant restriction on the choice of promises that results in simple, intuitive rules for policy conduct. These rules allow the Pareto gains from commitment to be attained, whilst simultaneously ensuring that policy choices will only vary in the underlying state of the economy. For instance, our chosen policy in the capital tax example allows tax rates to vary as the capital stock changes, but it does not allow different rates to be set for the same capital stock at different points in time. This is not true of Ramsey policy, as we confirm below.

Considering the specific examples in more detail, we show that the recursive Pareto principle justifies a constant, slightly positive inflation target in the basic linear-quadratic New Keynesian problem.[5] This contrasts with a Ramsey promise path in which inflation starts at a higher rate but gradually decays to zero. The loss associated with policy that satisfies the recursive Pareto criterion is initially above the loss from Ramsey, but after a finite length of time it becomes less appealing to continue with the Ramsey plan. We confirm an important feature of our method, which is generally true in deterministic examples without endogenous states: the constant policy that we identify also solves the problem of maximising the period-0 policymaker's objective when the choice set is restricted to

[4]This is known from the work of Marcet and Marimon (2015).

[5]Recall that the basic linear New Keynesian Phillips Curve takes the form:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \gamma y_t$$

where inflation is π_t, the output gap is y_t and β and γ are parameters. Since generically $\beta \in (0, 1)$, this exhibits a long-run trade-off between inflation and output. Thus permanently positive inflation is associated with a permanently positive output gap.

time-invariant policies only. This is not true of alternative methods for choosing time-invariant policy that have been proposed in the literature.

In the capital tax example we show that when promises satisfy the recursive Pareto criterion there is convergence to a steady state with mild positive capital taxes – around 28 per cent of net capital income in the calibration we use. Tax rates vary with the capital stock, but within relatively small bounds. Capital taxes are higher when the capital stock is high, because the marginal costs from distorting the accumulation decision are then relatively low. This contrasts with Ramsey policy, which requires initial capital tax rates of around 300 per cent of the tax base, gradually decaying to zero.[6] This result obtains even if the initial capital stock starts in its long-run steady-state level.[7] Analytically, the tax policy that we identify satisfies a set of intuitive restrictions linking measures of the labour wedge, consumption wedge and capital wedge in each time period. These restrictions are unchanging across time, and have close parallels with the classic 'inverse elasticity' rules that conventionally feature in static optimal tax theory.

The finding that steady-state taxes are positive when the recursive Pareto principle is satisfied is an important one because it emphasises that even if the celebrated 'zero capital tax' steady-state outcome does obtain under Ramsey policy, it cannot be viewed as a desirable outcome in isolation from the high-tax transition to it. Part of the reason why long-run capital taxes are held so low under Ramsey policy is to ensure that high initial capital levies will not discourage savings to too great an extent. When choice satisfies the recursive Pareto principle, high initial rates do not form part of the chosen plan. This reduces the need for early generations to be given incentives to save. Long-run capital taxes are higher as a consequence.

The model of social insurance with participation constraints provides a simple setting in which long-run outcomes are clearly undesirable under Ramsey policy. We assume that a fixed fraction of the population is permanently income-poor, and a utilitarian government would like to provide some degree of redistribution to these agents. Initially a Ramsey planner does this by redirecting to them some of the surplus raised by providing consumption insurance to the rest of the population, whose incomes are stochastic. Over time, however, the resources available for redistribution disappear. This is because initial high earners are promised high long-run utility levels, as compensation for the large payments they make to the scheme. Making good on these promises means it eventually becomes too costly to provide any further redistribution to low earners.

When promises are instead chosen to satisfy the recursive Pareto principle, redistribution to low earners is a feature of policy at every horizon. The chosen allocation is the best feasible stationary consumption distribution. Again, after a sufficient amount of time has passed it becomes preferable for all current and future policymakers to switch away from the continuation of Ramsey policy, to the policy that is recursively Pareto efficient.

In a final section before concluding, we highlight a close parallel between our recursive Pareto

[6] The tax base is net capital income, so tax rates in excess of 100 per cent imply that underlying capital assets are being taxed.

[7] Straub and Werning (2015) have recently shown that capital taxes may fail to converge to zero in the original version of the Judd (1985) model, with the economy instead collapsing to a corner solution under conventional parameterisations. Our version of the model differs from theirs in two regards. First, it is a representative-agent economy: there is no capitalist-worker distinction. Second, labour supply is endogenous. Together these features imply a zero steady-state capital tax rate for the standard calibrations we use.

approach to resolving Kydland and Prescott problems and the textbook analysis of Pareto-efficient dynamic allocations in pure consumption overlapping generations models, of the sort introduced by Samuelson (1958). Specifically, the problem that we solve when finding a recursively Pareto efficient sequence of promises appears isomorphic to the problem of finding a sequence of intergenerational transfers that will ensure long-run Pareto efficiency in the allocation problem among different generations in the OLG setting. The link between the two problems is helpful in clarifying aspects of our analysis, particularly the multiplicity issue studied in Section 6. We show that this is equivalent to the problem of fixing the initial consumption endowments of different generations in Samuelson's model. This parallel is likely to be useful in pointing to both normative and positive options for analysing how different generations should be treated.

1.1. Relation to literature

Many economists have expressed unease about the Ramsey approach to choice under time inconsistency. Svensson (1999) puts the problem succinctly, asking 'Why is period zero special?' The idea that preferences in some initial time period should be privileged above others does not seem consistent with the way that policy is, or should be, designed in practice. A large body of work has responded by supposing that choice must be made on a period-by-period basis, and analysing the resulting outcomes. An early branch of this literature, led by Chari and Kehoe (1990) and Atkeson (1991), investigated the set of reputational equilibria that could be supported in these settings by appropriate trigger strategies, which allowed for limited improvements on Markov-perfect equilibria.[8] More recently, a large number of papers have computed the properties of Markov-perfect equilibria, highlighting the inefficiencies that result. [9]

Our work shares with this literature an unease about the special treatment Ramsey policy gives to the first policymaker. Where it differs is in continuing to ask a normative question: How *should* policy be designed, given that we do not want period zero to be treated as special? This is different from asking what sorts of policies could be expected to follow if there were no formal commitment device.

When commitment devices are assumed, different branches of the literature have addressed perceived problems with 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. Influential work by Phelan (2006) and Farhi and Werning (2007) has placed particular focus on strategies for avoiding the immiseration result in the model of Atkeson and Lucas (1992). These authors proceed by attaching distinct non-zero Pareto weights to later generations when designing social insurance policy. Phelan considers policy that maximises steady-state welfare, whereas Farhi and Werning focus on a broader Pareto set. Both approaches are equivalent to increasing the social discount factor above its private-sector value, so that the government values later generations more than private individuals do. As an approach this is indeed sufficient to overturn immiseration, but its implications stretch far wider.

[8]Important papers by Sleet and Yeltekin (2006) and Golosov and Iovino (2014) have placed particular attention on the best equilibria that are supportable in this way.

[9]Examples include Klein and Ríos-Rull (2003), Ortigueira (2006), Ellison and Rankin (2007), Klein, Krusell and Ríos-Rull (2008), Díaz-Giménez et al. (2008), Martin (2009), Blake and Kirsanova (2012), Reis (2013) and Niemann et al. (2013).

It implies a change to dynamic decisionmaking even in models where no Kydland and Prescott problem is present. In a textbook Ramsey growth model, for instance, it would mandate a higher savings rate for any given capital stock.

An important message of our paper is that time inconsistency is a more limited problem than the broader question of how best to discount the future, and can be addressed via comparatively limited departures from standard dynamic optimisation. The main difference between our work and that of Phelan (2006) and Farhi and Werning (2007) is that we focus only on intergenerational preference differences with regards to the choice of promises. This is because, as we show, it is only the choice of promises that is subject to a time-inconsistency problem. Capital accumulation decisions will not be affected directly by what we do. More generally, our method will not mandate any departure from conventional policy choices in settings that do not feature Kydland and Prescott problems. This is not true when the social discount factor is changed. It follows from important differences between the Pareto criterion that we apply and the criterion that Farhi and Werning (2007) use. We defer a fuller discussion of it to Section 4.

The other body of work that has considered normative alternatives to Ramsey policy is the New Keynesian monetary policy literature. Here, 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.[10] 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.[11] But the justification for these approaches is 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? Our paper differs from this literature in the way it derives transition-free policy directly from the recursive application of an expanded choice set – the Pareto set, as distinct from the set of optimal choices – rather than by augmenting Ramsey policy *ex-post*.

1.2. Outline

The remainder of the paper proceeds as follows. Section 2 sets out the general problem that we study, and shows how this framework nests our three main examples. Section 3 explains the decomposition of the general problem into 'inner' and 'outer' components, and derives some important properties of the value function for the inner problem. Section 4 provides alternative Pareto criteria that could be applied the outer problem, and explains why we favour what we label an '*ex-post*' variant. Section 5 proves conditions under which a policy that satisfies the recursive Pareto principle will exist, and characterises the associated policy in steady state, highlighting its generic difference with Ramsey policy. Section 6 explores desirable selection procedures for obtaining a unique policy in the context

[10] 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, influenced by results due to Blake (2001) that showed timeless perspective policy may not maximise expected welfare.

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

of our three main examples. Section 7 highlights the parallels between recursively Pareto efficient policies and Pareto efficient equilibria in OLG economies. Section 8 concludes. All proofs are collected in an appendix.

2. General setup

Time is discrete, and runs from period 0 to infinity.[12] 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. ς 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 policymaker's preferences in period t are described by a time-separable objective criterion W_t:

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

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

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

a set of i contemporaneous restrictions linking controls and states:

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

and j infinite-horizon 'forward-looking' constraints:

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

It is assumed that $j \geq 1$ so the policymaker is subject to at least one forward-looking constraint, but otherwise there is no requirement that any of i, j or m should be non-zero. Constraints (2) to (4) 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) are allowed to vary directly in the vector of the exogenous stochastic process σ_s, as well as indirectly through $a_s(\sigma_s)$. Where the meaning is clear, we will usually keep the dependence of a_s on σ_s implicit, writing $h(a_s, \sigma_s)$ and so on.

The expectations in (4) 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 $P(\sigma'|\sigma)$ and a stationary probability of state σ denoted by $P(\sigma)$. Constraint (4) 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

[12] Notation is adapted from Marcet and Marimon (2015).

income risk. It will be useful at times to make reference to $A(\sigma) \subset \mathbb{R}^m$ as the space of control variables available for a given σ draw.

The exclusive focus on infinite-horizon constraints in (4) is solely in order to keep subsequent notation compact. The arguments that we make can easily be augmented to incorporate time-separable forward-looking restrictions of arbitrary order. Nonetheless, many constraints that do not immediately appear to have the same structure as (4) can be mapped into this form, as shown in the three examples below. The absence of state variables from the $h(\cdot)$ function is likewise for notational convenience. We have not experimented with examples that include such terms, but the extension ought to be straightforward.

It is constraints of the form (4) 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 satisfy them. 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 a quadratic sum of losses due to inflation and the output gap:

$$-\frac{1}{2}\sum_{s=t}^{\infty}\beta^{s-t}\left[\pi_s^2+\chi\left(y_s-\bar{y}\right)^2\right], \tag{5}$$

where $\chi > 0$ is a parameter and $\bar{y} > 0$ is the optimal level for the output gap.[13] The policymaker is subject to a single constraint, a deterministic version of the New Keynesian Phillips curve. This is usually written in the form:

$$\pi_s = \beta\mathbb{E}_s\pi_{s+1} + \gamma y_s, \tag{6}$$

with γ a parameter. Recursive substitution maps it into the same form as the general constraint (4):

$$\pi_s = \mathbb{E}_s\sum_{t=s}^{\infty}\beta^{t-s}\gamma y_s \tag{7}$$

where π_s corresponds to the $h_0(\cdot)$ function and γy_t corresponds to $h(\cdot)$.[14] 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).

[13] In the underlying non-linear version of the problem, this exceeds the natural level of output due to monopoly power in the product market.

[14] The restriction is written as an equality rather than the inequality treated in the general setting. This is not a significant departure: it would always be possible to write the restriction less compactly as two overlapping inequalities.

Example 2: A capital tax problem

We consider 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:

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

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):

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

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 examples of this type this is usually written in the form:

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

Again, recursive substitution allows it to be mapped into the structure of our general constraint:

$$\sum_{t=s}^{\infty} \beta^{t-s}\left[u_{c,t+1} c_{t+1} + u_{l,t+1} l_{t+1}\right] \geq u_{c,s}\left(c_s + k_{s+1}\right) + u_{l,s} l_s \tag{12}$$

so that $u_{c,t}\left(c_t + k_{t+1}\right) + u_{l,t} l_t$ corresponds to $h_0\left(\cdot\right)$ and $u_{c,t+1} c_{t+1} + u_{l,t+1} l_{t+1}$ corresponds to $h\left(\cdot\right)$. This constraint says that the value of consumption and capital purchases in period $s+1$, net of any labour income that period, must weakly exceed the value of the capital holdings that are taken into period $s+1$, where these values are calculated in period s 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.[16] This will conflict with the incentive of a later policymaker to tax the consumer's existing, inelastic capital holdings.

[15] See Chari and Kehoe (1999).

[16] Stating the condition as an inequality implicitly allows that the policymaker could make use of positive lump-sum transfers in any given time period, but taxes must be distortionary.

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, and to keep the discussion compact we restrict attention 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:

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

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:

$$\sum_{s=t}^{\infty} \beta^{s-t} \left[(1-\mu) \sum_{\sigma=0}^{\infty} (1-p)^{\sigma} pu\left(c_s(\sigma)\right) + \mu u(c_s^p) \right], \tag{13}$$

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:

$$(1-\mu) \sum_{\sigma=0}^{\infty} (1-p)^{\sigma} pc_s(\sigma) + \mu c_s^p \leq [1-(1-\mu)p] y^l + (1-\mu) py^h. \tag{14}$$

Incentive compatibility requires the insurance scheme to deliver at least as much utility to an agent as they could obtain in autarky. For agents with stochastic income draw y_t, this implies infinite-

[17] This keeps 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.

[18] We assume 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.

horizon forward-looking constraints of the form in (4):

$$\mathbb{E}_s \sum_{t=s}^{\infty} \beta^{t-s} u\left(c_t\left(\sigma_{t,i}\right)\right) \geq u\left(y_s\right) + \frac{\beta}{1-\beta}\left[pu(y^h) + (1-p)\, u(y^l)\right], \tag{15}$$

so that the autarky value on the right-hand side corresponds to $h_0\left(\cdot\right)$, and the within-period utility function $u\left(c_t\left(\sigma\right)\right)$ corresponds to $h\left(\cdot\right)$. For the permanently low-income agents, the constraint is:

$$\sum_{s=t}^{\infty} \beta^{s-t} u(c_t^p) \geq \frac{1}{1-\beta} u(y^l) \tag{16}$$

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, and h_0 later in the analysis.

Assumption 1. *The functions $r : X \times A \to \mathbb{R}$, $l : X \times A \to \mathbb{R}^n$, $p : X \times A \to \mathbb{R}^i$, $h : A\left(\sigma\right) \to \mathbb{R}^j$ and $h_0 : X \times X \times A\left(\sigma\right) \to \mathbb{R}^j$ 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, and h_0 are continuously differentiable.*

Assumption 3. *The function h is quasi-concave and the function h_0 is quasi-convex.*

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

Assumption 5. *The function r is strictly quasi-concave.*

Assumption 6. *The function r is strictly concave.*

Assumption 7. *The functions p and h are concave and the function h_0 is convex.*

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 to 7 place restrictions of varying strength on the structure of the problems studied. In the capital tax example, for instance, the relevant functions in the implementability constraint (12) will not generally satisfy quasi-concavity, let alone full concavity. For an example of this type it will generally only be possible to confirm Assumptions 3 to 7 under specific functional forms.

The main purpose of Assumptions 3, 5 and 7 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

Before setting out our recursive Pareto approach to choosing policy, it is useful briefly to recap the properties of Ramsey policy, both in general and in our three examples.

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 (4) 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.

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 to 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.

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

[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.

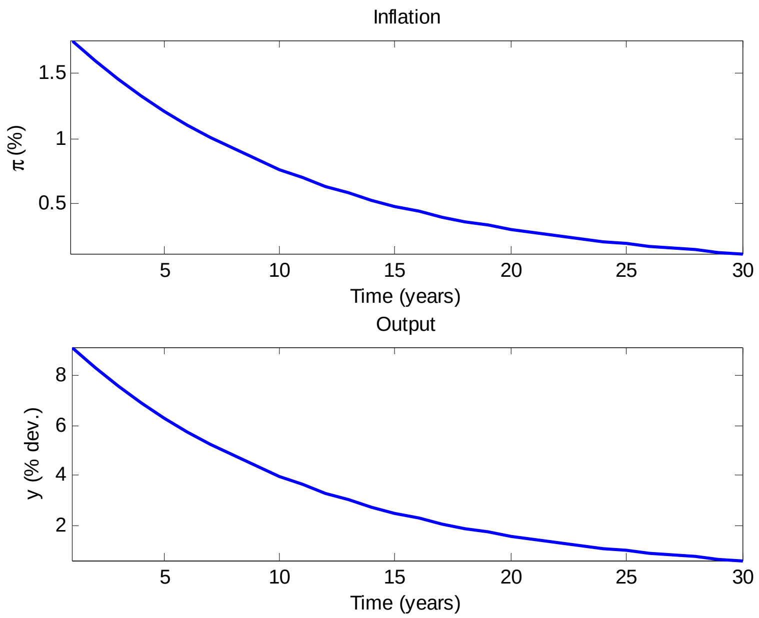

Figure 1: Ramsey paths for inflation and output in the linear-quadratic inflation bias problem

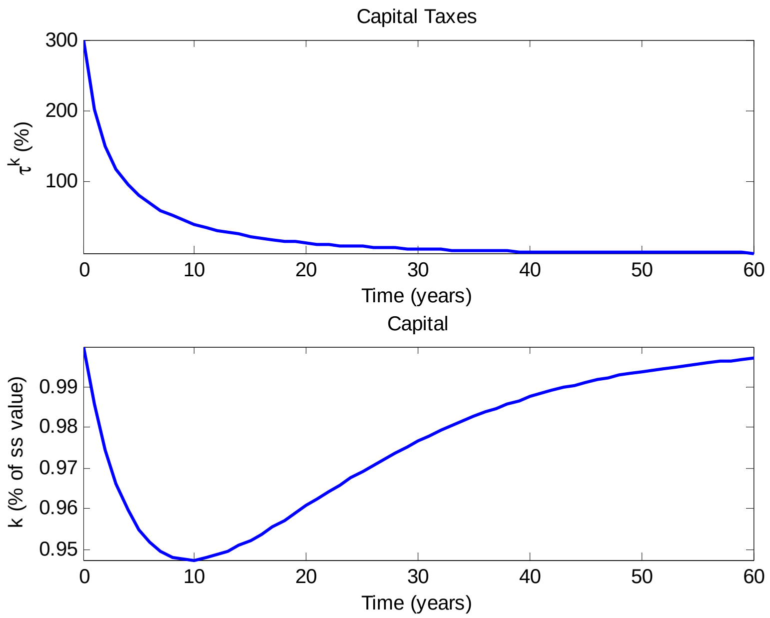

Figure 2: Ramsey capital taxes and the capital stock in the capital tax problem

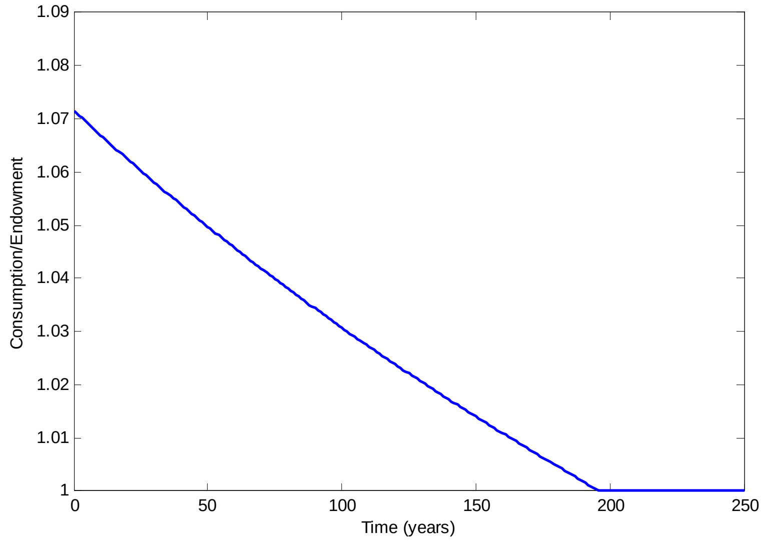

Figure 3: Consumption path for low-income agents in the social insurance problem

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 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 initial policy with capital taxes at 300 per cent could hardly be more different from the limiting policy when capital taxes converge on zero.

[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.

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

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.

This asymmetry between long-run and short-run outcomes has led a number of authors to investigate whether an alternative approach to policy design could be possible. The question is what the relevant 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 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 strategy. 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 expand the within-period choice set, defining a policy criterion that is weaker than 'optimality' but can still be applied recursively. Specifically, we focus on a version of the Pareto criterion applied to the choice of promises over time. In the next section we specify exactly

[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 (2015) have recently highlighted the inseparability of transition dynamics from the optimality of the subsequent zero rate.

[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 (2015) that the cross-sectional Pareto weights applied by the policymaker to agents will evolve in a non-stationary manner in environments such as our Example 3. Woodford (2003) presents his approach as imposing steady-state values for the multipliers on past promises when solving the period 0 decision problem. But in Example 3 these multipliers are one and the same as the cross-sectional Pareto weights. They do not have steady-state values because of their non-stationarity. It is not clear how a timeless perspective policymaker should proceed.

what we mean by 'promises' in this context, and how these objects differ from other choice variables.

3. Inner and outer problems

Our approach is to divide the general problem 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 we study them as complete dynamic sequences. We consider the choice of these sequences to be an 'institutional design' problem. An inner problem can then be cast, determining the optimal choice of all other variables of interest, holding constant a given sequence of promises. 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.

3.1. The inner problem

We present the inner problem first. This 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:

$$h\left(a_s, \sigma_s\right) + \beta \mathbb{E}\omega_{s+1}\left(\sigma_{s+1}\right) \geq h_0\left(a_s, x_s, x_{s+1}, \sigma_s\right), \tag{17}$$

$$h\left(a_s, \sigma_s\right) + \beta \mathbb{E}\omega_{s+1}\left(\sigma_{s+1}\right) \geq \omega_s\left(\sigma_s\right), \tag{18}$$

for all $s \geq t$, where $x_t \in X$ is the initial state vector and $\omega_s\left(\sigma_s\right) \in \mathbb{R}^j$ is a σ-contingent vector of promise values in all $s \geq t$.[30]

The collection of promise values across σ_s draws is denoted by $\omega_s := \{\omega_s\left(\sigma_s\right)\}_{\sigma_s \in \Sigma}$. Condition (17) can be interpreted as a 'promise-making' constraint, as it represents a new cross-restriction on the choice of policy variables in s, conditional on a particular array of promise values for period $s+1$. Condition (18) is a 'promise-keeping' constraint, as it ensures that prior commitments, in the form of $\omega_s\left(\sigma_s\right)$, are met.

For the original constraint (4) to hold for all $s \geq t$, it is sufficient to ensure that (17) holds for all $s \geq t$ and (18) for all $s > t$. Imposing (18) 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

[29]The separation most closely resembles that of Hansen and Sargent (2008) in their specification of robust control theory.
[30]When (17) binds we additionally require for consistency that $\omega_s\left(\sigma_s\right) = h_0\left(a_s, x_s, x_{s+1}, \sigma_s\right)$.

restrictions that feature from period t onwards. Clearly this is a manifestation of time inconsistency.

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 $\Omega(x_t) \subset (\mathbb{R}^j \times \mathbb{R}^\varsigma)^\infty$, for an initial state vector x_t. The interior of $\Omega(x_t)$ is represented by $\mathring{\Omega}(x_t)$. Compactness of A and X and the continuity properties imposed under Assumption 1 together imply that the constraint functions in (17) and (18) are bounded uniformly in s. This in turn means that sufficiently large, uniform bounds can be imposed on the value of $\omega_s(\sigma_s)$ without affecting the problem. These bounds are incorporated into the definition of $\Omega(x_t)$:[31]

Assumption. $\{\omega_s\}_{s=t}^{\infty} \in \Omega(x_t)$ *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 $\{\omega_s\}_{s=t}^{\infty} \in \Omega(x_t)$. This object will be a major focus of the analysis, and we label it the 'promise-value function'. Where 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 $\left(\mathbb{R}^j \times \mathbb{R}^\varsigma\right)^\infty$.

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. The following property follows from a standard application of the maximum theorem to this setting. A proof is given in the appendix:

Proposition 1. *Suppose Assumptions 1, 3, 4 and 5 hold. Fix x_t and $\{\omega_s\}_{s=t}^{\infty} \in \mathring{\Omega}(x_t)$. For all $T \geq t$, the sequence of state vectors $\{x_{s+1}^*\}_{s=t}^{T}$ induced by $\left\{x_t, \{\omega_s\}_{s=t}^{\infty}\right\}$ is **continuous** in $\{\omega_s\}_{s=t}^{\infty}$ and x_t.*

The qualification that promises should be in the interior of $\Omega(x_t)$ here simply reflects the fact that outside of $\Omega(x_t)$ the optimal path $\{x_{s+1}^*\}_{s=t}^{\infty}$ is not defined, so continuity cannot be satisfied.

3.1.2. Time consistency

Proposition 2. *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. It follows because the inner problem has seen all forward-looking restrictions replaced by a succession of separate static constraints, (17) and (18). Static constraints cause no time inconsistency issues. The important insight here 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.

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

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 plays a central role in the arguments that follow, and the strength of these arguments will often depend in turn upon the regularity properties of V. In this subsection we thus set out in detail the implications for V of placing different combinations of Assumptions 1 to 7 on the problem's primitives. Proofs rely on established properties of parameterised optimisation problems, and are contained in the appendix. A reader less interested in these technical details can skip to Section 3.3 without losing track of the main arguments.

Proposition 3. *Fix x_t and $\{\omega_s\}_{s=t}^{\infty} \in \mathring{\Omega}(x_t)$. Suppose Assumptions 1 and 3 hold. Then the promise-value function $V\left(\{\omega_s\}_{s=t}^{\infty}, x_t\right)$ is **continuous** in $\{\omega_s\}_{s=t}^{\infty}$. Suppose Assumptions 1 and 4 hold. Then the promise-value function $V\left(\{\omega_s\}_{s=t}^{\infty}, x_t\right)$ is **continuous** in x_t.*

The proof 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 in most settings to strengthen it to continuous differentiability. This requires a standard constraint qualification condition, LICQ, to be satisfied by the restrictions (17) and (18) at a chosen allocation. This condition is presented in the appendix.[32] It amounts to requiring that each binding constraint in (17) and (18) is affected in a linearly independent manner by changes in the policy variables, and thus that there is a unique set of Lagrange multipliers associated with the inner problem. This does not seem an important limitation for practical purposes.

Proposition 4. *Suppose Assumptions 1 to 5 hold. 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 $j \times 1$ vector:*

$$-\beta^{s-t} P(\sigma_s)\, \lambda_s^k(\sigma_s) + \beta^{s-t} P(\sigma_s) \left[\lambda_{s-1}^m(\sigma_{s-1}) + \lambda_{s-1}^k(\sigma_{s-1})\right] \tag{19}$$

where $\lambda_s^m(\sigma_s)$ and $\lambda_s^k(\sigma_s)$ are the vector multipliers on constraints (17) and (18) respectively,[33] and:

$$\lambda_{t-1}^m(\sigma_{t-1}) = \lambda_{t-1}^k(\sigma_{t-1}) = 0$$

The main contribution of Proposition 4 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 (19) 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

[32] It is based on Wachsmuth (2013).

[33] The superscripts distinguish the multipliers on promise-making and promise-keeping constraints respectively.

promises is only a cost, given by the first term in (19). 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 (17) and and (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 (2015). It has important implications for the character of the Ramsey solution in the long run, particularly in models with dynamic incentive constraints.

Proposition 5. *Suppose Assumptions 1, 5 and 7 hold. Fix x_t. The promise-value function $V(\cdot, x_t)$ is **strictly quasi-concave** in $\{\omega_s\}_{s=t}^{\infty} \in \Omega(x_t)$ so long as the corresponding constraints of the inner problem bind.*[34] *If Assumptions 1, 5 and 7 hold except that r is only weakly concave then $V(\cdot, x_t)$ is **quasi-concave** in $\{\omega_s\}_{s=t}^{\infty} \in \Omega(x_t)$ where the corresponding constraints of the inner problem bind.*

Proposition 6. *Suppose Assumptions 1, 5 and 7 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 6 relates solely to the constraint set, so goes through without any concavity restrictions on r. For this reason we only prove 5 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 exclusively 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 $\left(\mathbb{R}^j \times \mathbb{R}^\varsigma\right)^\infty$. 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 $\{\omega''_s\}_{s=0}^{\infty} \in \left(\mathbb{R}^j \times \mathbb{R}^\varsigma\right)^\infty$ 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:

[34] That is, for variations in ω_t that have an impact on binding constraints of the form (17) and (18)

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 $\{\omega'_s\}_{s=0}^{\infty} \in \left(\mathbb{R}^j \times \mathbb{R}^{\varsigma}\right)^{\infty}$ such that $V(\{\omega_s'\}_{s=0}^{\infty}, x_0) > V(\{\omega_s^R\}_{s=0}^{\infty}, x_0)$.

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

Proposition 7. *Let $\{\omega_s^R\}_{s=0}^{\infty}$ be a Ramsey plan, given some x_0. Exactly one of the following is true:*

1. *Constraints (17) and (18) 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. 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 4 means a necessary optimality condition with respect to the choice of $\omega_s(\sigma_s)$ is:

$$\lambda_s^k(\sigma_s) = \lambda_{s-1}^m(\sigma_{s-1}) + \lambda_{s-1}^k(\sigma_{s-1}), \tag{20}$$

for $s > 0$, and:

$$\lambda_0^k(\sigma_0) = 0 \tag{21}$$

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 (2015). 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 principally by the need to make good on past promises, rather than the priorities contained in the underlying social welfare objective W_t.

4. The Pareto approach to designing promises

This section explains our recursive Pareto criterion for choosing promise sequences. The benefit of analysing time inconsistency through the apparatus of the promise-value function is that $V(\{\omega_s\}_{s=t}^{\infty}, x_t)$ can be treated as a standard preference ordering over promise sequences from period t onwards. Different generations of policymakers have different preferences regarding the overall promise sequence that is chosen, just as different agents in a market economy will generally disagree about the overall allocation of social resources. But this does not mean that Pareto gains cannot be found.[35]

4.1. Pareto efficiency: definition

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. An ideal choice of promises would be a time-consistent optimal choice. That is, a sequence $\{\omega_s\}_{s=0}^{\infty}$ with the property that there exists no alternative sequence $\{\omega_s'\}_{s=0}^{\infty}$ such that the policymaker in *some* period t would prefer to switch to the continuation $\{\omega_s'\}_{s=t}^{\infty}$ rather than sticking with $\{\omega_s\}_{s=0}^{\infty}$, given an inherited state vector x_t induced by $\{\omega_t\}_{s=0}^{\infty}$. As Proposition 7 has confirmed, this is not possible. A looser desirability requirement is that there should not exist an alternative sequence $\{\omega_s'\}_{s=0}^{\infty}$ such that the policymaker in *every* period $t \geq 0$ would prefer to switch to the continuation $\{\omega_s'\}_{s=t}^{\infty}$ rather than sticking with $\{\omega_s\}_{s=t}^{\infty}$, again given an inherited state vector induced by $\{\omega_s\}_{s=0}^{\infty}$. This is a far less demanding restriction to place on choice, and one that seems uncontroversial. It is hard to justify an institution that every generation would prefer to be rid of.

This verbal definition has considered Pareto efficiency with respect to the preference orderings of every policymaker from period 0 onwards. Ultimately we are interested in a desirability criterion that can be applied recursively. For this reason the formal definition does not fix the initial date at zero:

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

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

for some scalar$\varepsilon > 0$, independent of τ.

The particular definition here is 'weak' Pareto criterion, in the sense that we include in the Pareto set every promise sequence that is not *strictly* dominated by another for all policymakers.[36] This makes the Pareto set as large as possible, and this in turn is necessary if we are to find a promise sequence that satisfies the above definition recursively.[37] Note that the Ramsey promise sequence $\{\omega_s^R\}_{s=0}^{\infty}$ is Pareto efficient from period 0. It will never be possible to find an alternative sequence that the policymaker in period 0 would strictly prefer. But it does not follow that the continuation of the Ramsey sequence from $t > 0$, $\{\omega_s^R\}_{s=t}^{\infty}$ will be Pareto efficient from t.

[35] In Section 7 we make the parallel with a market economy explicit, showing that the outcome of our choice criterion can be decentralised by a Walrasian mechanism that allows each generation of policymaker to trade the promises it makes and the promises it keeps.

[36] 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 still further the set of $\{\omega_s\}_{s=t}^{\infty}$ that can be included.

[37] Heuristically, a promise sequence $\{\omega_s^*\}_{s=t}^{\infty}$ such that ω_t^* places a binding constraint on the policymaker at t is always weakly dominated from t onwards by one that neglects this initial commitment. This means a Pareto efficiency concept that is based on weak dominance could never be applied recursively whilst allowing for binding commitments to be made. As we note in Section 7, this is a counterpart to the common observation in two-period overlapping generations models that the core is empty: every young generation and its successors can improve on a transfer scheme that allocates resources to the current old.

Discussion

An important feature of this definition is that it is applied to the 'on-path' preferences of policymakers. That is, it takes as given that the policymaker in $t+\tau$ has inherited a state vector consistent with the promise sequence $\{\omega_s^*\}_{s=t}^{\infty}$, and asks whether they regret continued commitment to this sequence. An alternative possibility would be to compare sequences $\{\omega_s^*\}_{s=t}^{\infty}$ and $\{\omega_s'\}_{s=t}^{\infty}$ under the assumption that choice between the two is made 'once and for all' in period t, and so the value function for the policymaker at $t+\tau$ when $\{\omega_s'\}_{s=t}^{\infty}$ is being considered includes some state vector $x_{t+\tau}'$, induced by $\{\omega_s'\}_{s=t}^{\infty}$, instead of $x_{t+\tau}^*$. This sort of '*ex-ante*' approach to Pareto efficiency is similar to the definition used by Farhi and Werning (2007, 2010), except that they apply it not just to the choice of promises (and induced states) but to the entire allocation, including the variables that we include in the inner problem.

It is possible to make sensible normative arguments for both sorts of definition. One advantage of ours is that it depends on preferences that can actually be observed along the equilibrium path, rather than on an *ex-ante* assessment of what future generations will want. As we show in Section 7, it is the efficiency concept that is satisfied when policymakers from different generations are allowed to trade the promises that they make and keep over time, in a Walrasian market setting.

More significantly for our present purposes, the definition that we have given preserves the exclusive focus on choices that are subject to a time-inconsistency problem. If we did not treat the inherited state vector as given for each generation, Pareto comparisons could be influenced by the fact that later generations would prefer earlier generations to accumulate more capital, for instance. This is an important normative consideration, but the standard practice in the policy design literature is to leave control over variables set in t to policymakers alive in t or earlier.[38] It is only when these 'pre-t' generations disagree that a time-inconsistency *problem* arises. Our definition focuses on Pareto efficient resolutions to this disagreement alone. By contrast, Farhi and Werning (2007, 2010) consider the Pareto frontier that is traced out as the social discount factor is varied. This has implications far beyond Kydland and Prescott settings.

4.2. Recursive Pareto efficiency

Once the definition of Pareto efficiency is established, recursive Pareto efficiency follows straightforwardly:

Definition. A promise sequence $\{\omega_s^*\}_{s=0}^{\infty}$ is **recursively Pareto efficient** (RPE) if and only if the continuation sequence $\{\omega_s^*\}_{s=t}^{\infty}$ is Pareto efficient from all periods $t \geq 0$.

Heuristically, recursive Pareto efficiency requires not just that a dynamic promise sequence should not be regretted by all current and future policymakers when it is first implemented, but also that it should not *come* to be uniformly regretted in this manner later – once any benefits to initial generations have passed. There is no guarantee that Ramsey policy will satisfy this requirement. Ramsey policy is Pareto efficient in period 0, but its continuation may not be thereafter.

[38]Of course, this does not stop these policymakers having an altruistic concern for future outcomes.

5. Recursively Pareto-efficient policies

This section explores the implications for policy of adopting a recursive Pareto criterion. In particular, it provides conditions under which the recursive Pareto criterion can generically be satisfied, and provides a general characterisation of the steady-state properties of the associated policy. This differs systematically from Ramsey policy.

5.1. Recursive Pareto efficiency: existence

In this subsection we prove a set of conditions under which the Pareto criterion can be satisfied recursively. The necessary conditions for the result to go through are stated in terms of the properties of the promise-value function and induced choice for the endogenous state variables. Sufficient conditions on the primitives for these properties to hold are, in turn, stated in Propositions 1, 3, 5 and 6. The main Proposition is as follows:

Proposition 8. *Suppose the following conditions hold:*

1. *The promise-value function $V\left(\{\omega_s\}_{s=0}^{\infty}; x_0\right)$ is strictly quasi-concave in $\{\omega_s\}_{s=0}^{\infty}$ (or invariant where the corresponding constraints do not bind).*
2. *The optimal choice of state variables for the inner problem, $\{x_s^*\}_{s=1}^{T}$, is continuous in all elements of $\{\omega_s\}_{s=0}^{\infty}$ for all $T \geq 0$.*
3. *$V\left(\{\omega_s\}_{s=0}^{\infty}; x_0\right)$ is continuous in x_0.*
4. *$\Omega\left(x_0\right)$ is non-empty and convex.*

Then there exists a sequence of promises $\{\omega_s\}_{s=0}^{\infty} \in \Omega\left(x_0\right)$ that satisfies recursive weak Pareto efficiency.

The proof of this Proposition, contained in the appendix, relies on a non-trivial fixed-point argument. It defines a correspondence from $\Omega\left(x_0\right)$ to itself with the property that this correspondence identifies strict Pareto improvements on any given $\{\omega_s\}_{s=t}^{\infty}$, from the perspective of some $t \geq 0$. A fixed point of the correspondence, whose existence follows from Kakutani's theorem, must be a promise sequence that satisfies the weak Pareto criterion recursively.

Conditions 1 to 4 in the Proposition should be viewed as sufficiency requirements for existence: they are not necessary. In many interesting examples they may not be satisfied, but this does not rule out the possibility that the recursive Pareto criterion could be satisfied. This includes the dynamic capital tax problem that we consider as our Example 2.

5.2. Recursive Pareto efficiency: long-run characterisation

This section characterises the steady-state outcomes that obtain when RPE is used to design policy. These outcomes will be shown to differ systematically from steady-state outcomes under Ramsey policy. The main general characterisation result is expressed in terms of the multipliers on the promise constraints (17) and (18):

Proposition 9. *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*. If this promise sequence satisfies RPE then it is not possible to find a time period* $\tau \geq 0$ *and a scalar* $\varepsilon > 0$ *such that either:*

$$\lambda_s^k(\sigma') - \beta \sum_{\sigma \in \Sigma} \frac{P(\sigma'|\sigma) P(\sigma)}{P(\sigma')} \left[\lambda_s^m(\sigma) + \lambda_s^k(\sigma)\right] \geq \varepsilon \tag{22}$$

for all periods $s \geq \tau$*, or:*

$$\lambda_s^k(\sigma') - \beta \sum_{\sigma \in \Sigma} \frac{P(\sigma'|\sigma) P(\sigma)}{P(\sigma')} \left[\lambda_s^m(\sigma) + \lambda_s^k(\sigma)\right] \leq -\varepsilon \tag{23}$$

for all periods $s \geq \tau$*.*

The measure $\frac{P(\sigma'|\sigma)P(\sigma)}{P(\sigma')}$ here is a 'reverse' transition probability that the predecessor to an observed state σ' was state σ. It satisfies:

$$\sum_{\sigma \in \Sigma} \frac{P(\sigma'|\sigma) P(\sigma)}{P(\sigma')} = 1. \tag{24}$$

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 and the contemporary participation constraint generally binds. In this case lagged multipliers do not affect the allocation.[39] The linear-quadratic inflation bias and capital tax examples are both fully deterministic, so for these the probabilistic term can be dropped.

The main implication of this Proposition is that the recursive Pareto criterion places limits on the long-run trade-off that is struck between the benefits from making promises and the costs of keeping them. The first multiplier terms on the left-hand-side of inequalities (22) and (23) are the shadow costs in period s of increasing the promise that is 'kept' in state σ in that period, $\omega_s(\sigma)$, by a unit. The second multiplier terms capture the shadow benefits from simultaneously increasing the promise that is 'made' in period s for the same state in $s+1$, i.e. $\omega_{s+1}(\sigma)$ – also by a unit. The discount factor β captures the fact that this increase will occur one period later. For a promise sequence to be recursively Pareto efficient, uniform increases or decreases in promises in all time periods cannot be to the benefit of all policymakers. The inequalities ensure that such uniform improvements are not possible.

This result has important implications for the character of RPE policy in steady state, under the assumption that such a state exists. We state these formally as a Corollary.

Corollary. *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 (17) and (18) converge to* $\lambda_{ss}^m(\sigma)$ *and* $\lambda_{ss}^k(\sigma)$ *respectively. Then:*

[39] That is, the value of $\lambda_{ss}^m(0) + \lambda_{ss}^k(0)$ is independent of the value of $\lambda_{ss}^k(0)$.

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

$$\lambda_{ss}^{k}\left(\sigma'\right)=\beta\sum_{\sigma\in\Sigma}\frac{P\left(\sigma'|\sigma\right)P\left(\sigma\right)}{P\left(\sigma'\right)}\left[\lambda_{ss}^{m}\left(\sigma\right)+\lambda_{ss}^{k}\left(\sigma\right)\right]. \tag{25}$$

2. *If the promise-value function $V\left(\cdot,x\right)$ is quasi-concave for all $x\in X$, then the promise sequence satisfies RWPE whenever (25) holds.*

The results here are analogous to a standard first-order condition. They provide a necessary condition for RPE to hold, which becomes sufficient when the relevant objectives are quasi-concave.[40]

In all three examples, condition (25) 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, β. This feature is notable, because it contrasts with the outcome under Ramsey policy, characterised by (20):

$$\lambda_{s}^{k}\left(\sigma_{s}\right)=\left[\lambda_{s-1}^{m}\left(\sigma_{s-1}\right)+\lambda_{s-1}^{k}\left(\sigma_{s-1}\right)\right] \tag{26}$$

This expression implies a multiplier recursion similar to (25), but without any downward drift. Instead, the coefficient on past promise-making constraints under Ramsey-optimal policy is fixed to one. Intuitively, the Ramsey policymaker trades off the benefits today of making a promise for tomorrow with the costs tomorrow of keeping that promise. Under RPE policy, the enduring cost today of keeping yesterday's promise is assessed simultaneously to the benefit from making an improved promise for tomorrow. This implies a crucial difference in timing: keeping a promise today is more onerous than keeping one tomorrow, so long as $\beta < 1$.

In models with infinite-horizon constraints of the form (4), the multipliers are non-stationary, and often will not converge to a steady state. This in turn can impart uncomfortable long-run properties, such as the absence of long-run redistribution in the social insurance example. It is also the driving force behind the immiseration result in dynamic models of asymmetric information, and the conclusion of Straub and Werning (2015) that Ramsey-optimal capital taxes in the Judd (1985) framework can imply 'corner' outcomes, with zero long-run consumption for workers. When RPE policy is instead considered, the downward drift in the multipliers generally allows steady state to exist, and the continued satisfaction of the Pareto criterion ensures that its properties are more benign than these extreme cases.

In many other models of interest, Ramsey policies *do* converge to a steady state with constant multipliers. This is true, for instance, in both the inflation bias example and the capital tax example that we have outlined above. In these cases it is clear that the steady-state version of (26) will directly contradict (25). Thus the steady state of Ramsey policy does not satisfy the Pareto criterion. We interpret this as implying that the steady state of Ramsey policy cannot be justified in isolation from transition dynamics. To take the most well-known example, taxing capital income is only a bad idea after the initial high-tax Ramsey trajectory.

[40] Though the intuition behand Part 2 of the Corollary is standard, the dynamic horizon causes some complications. We provide a proof in the Appendix for completeness.

Proposition 9 relates solely to long-run outcomes. It does not specify any properties of the transition dynamics associated with RPE promise sequences. Indeed, if the value function is quasi-concave then it follows from the Corollary that if one promise sequence satisfies RPE 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 RPE.[41] This leaves us with significant indeterminacy in the solution concept as presented so far. In the next section we consider how to resolve this.

6. Transition dynamics and policy rules `[under revision]`

The results of Section 5 are important because they provide a justification for policies that converge to outcomes that are different from the long-run outcomes of Ramsey policy. This matters because the long-run outcomes under Ramsey policy have often been treated as desirable *per se* – an idea formalised in the 'timeless perspective' approach of Woodford (2003). Yet in order to use our approach to devise practical rules for policymaking we need to address the fact that many transition paths may be consistent with RPE policy. This section proposes one way to do this, resulting in easily interpretable policy rules across the different examples that we consider. It exploits the additive separability of the objective and constraint functions over time, to ensure that policy choices over control variables, a_t, will be time-invariant functions of t-dated objects alone. The aim is to translate the general theory into practice, and for this reason we find it simplest to explain our approach by direct reference to our chosen three examples. We consider them in turn.

Example 1: Linear-quadratic inflation bias

The inner problem for this example, viewed in period 0, is the following:

$$\max_{\{\pi_t, y_t\}_{t=0}^{\infty}} -\frac{1}{2}\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2+\chi\left(y_t-\bar{y}\right)^2\right]$$

subject to:

$$\pi_t = \gamma y_t + \beta\omega_{t+1} \tag{27}$$
$$\omega_t = \gamma y_t + \beta\omega_{t+1} \tag{28}$$

and given $\{\omega_t\}_{t=0}^{\infty}$. Necessary conditions for an optimum are:

$$-\pi_t + \lambda_t^m = 0 \tag{29}$$
$$-\chi\left(y_t-\bar{y}\right) - \gamma\left(\lambda_t^m+\lambda_t^k\right) = 0 \tag{30}$$

where λ_t^m and λ_t^k are the multipliers on the promise-making and promise-keeping constraints (27) and (28) respectively. We know that if the chosen policy is to satisfy recursive Pareto efficiency, the

[41] This property will in fact hold regardless of quasi-concavity, but quasi-concavity is sufficient to confirm RPE once a steady-state satisfying (25) is found.

only possible steady state is such that $\beta\left(\lambda_{ss}^{m}+\lambda_{ss}^{k}\right)=\lambda_{ss}^{k}$. This is the implication in the current setting of the Corollary to Proposition 9.

The question is how to choose the complete $\{\omega_t\}$ sequence – or, equivalently, how to restrict the relative sizes of the multipliers λ_t^m and λ_t^k. Since the setting is a purely stationary one, it is clear that the most appealing choice will be one that delivers similar stationarity in the resulting policy rule. One of our two main motivations for departing from Ramsey policy was that time-varying policies for a stationary economic system seemed contrary to the demands of practical institutional design. For this reason we take the limiting restriction on the multipliers that is necessary according to Proposition 9, and impose it every time period. That is, we set:

$$\beta\left(\lambda_t^m+\lambda_t^k\right)=\lambda_t^k \tag{31}$$

for all t. Notice that this is a within-period restriction linking the shadow cost of keeping promises in period t to the shadow benefit of making promises for $t+1$. This contrasts with the equivalent Ramsey restriction, which would equate the shadow costs and benefits associated with changing a single promise:

$$\lambda_t^k=\left(\lambda_{t-1}^m+\lambda_{t-1}^k\right)$$

Whereas the Ramsey policymaker optimises distinctly over each promise, the RPE policymaker in period t is behaving as if they face a linear 'pricing' restriction linking the promise that they can make for tomorrow to the promise that they keep today. There is a very close link here with the role of intergenerational transfers in achieving Pareto improvements in textbook overlapping generations models of the sort introduced by Samuelson (1958). We will formalise this intuition more fully in Section 7 below.

Condition (31) allows the optimality conditions from the inner problem to be condensed to a single, time-invariant restriction:

$$\pi_t=(1-\beta)\frac{\chi}{\gamma}\left(\bar{y}-y_t\right) \tag{32}$$

When substituted into the original New Keynesian Phillips Curve (6), this gives a first-order difference equation in y_t. This equation has a unique non-explosive solution, in which:

$$\begin{aligned} y_t &= \frac{(1-\beta)^2\chi}{(1-\beta)^2\chi+\gamma^2}\bar{y} \\ \pi_t &= \frac{(1-\beta)\chi\gamma}{(1-\beta)^2\chi+\gamma^2}\bar{y} \end{aligned}$$

for all t. It is straightforward to check that this is also the solution to the simple problem of finding values for y_t and π_t that maximise the policy objective on the set of entirely constant policies. This provides a useful endorsement of our approach. In a purely stationary setting such as this, a policy that does not have the transition dynamics associated with Ramsey must be a constant policy: there are no state accumulation dynamics to complicate matters. The fact that RPE policy is able to select the best constant policy contrasts with the 'timeless' approach of selecting the Ramsey steady-state policy. Figure 4 compares the paths of inflation and output under Ramsey and stationary RPE pol-

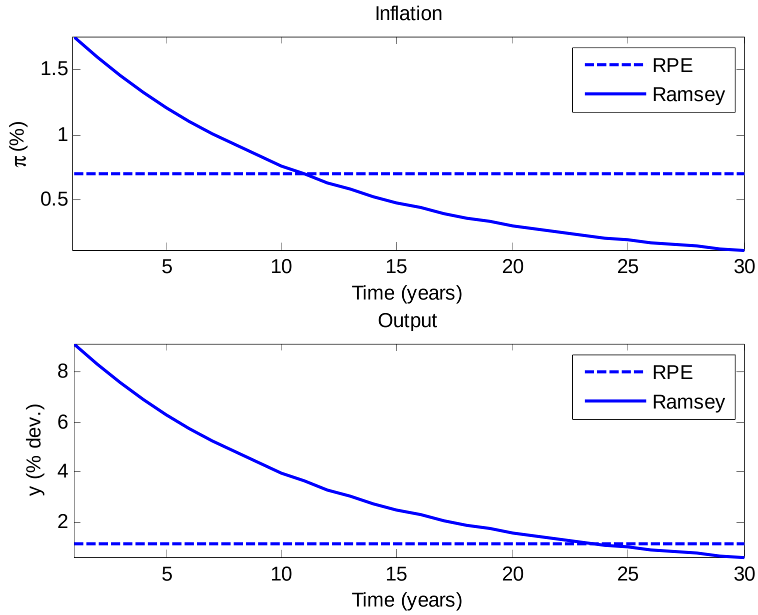

Figure 4: Ramsey and stationary RPE policy in the inflation bias example

icy.[42] There is a visible sense in which RPE policy delivers a convex combination of the short-run and long-run outcomes from Ramsey. It has inflation permanently above zero, which permits a small positive output gap to exist in perpetuity.[43]

Example 2: Capital taxation

The inner problem for this example can be written as follows:

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

subject to:

$$\begin{aligned} c_t + k_t + g &\leq F\left(k_t, l_t\right) + (1-\delta) k_t && (33) \\ u_{c,t} c_t - v_{l,t} l_t + u_{c,t} k_{t+1} &\leq u_{c,t} c_t - v_{l,t} l_t + \beta \omega_{t+1} && (34) \\ \omega_t &\leq u_{c,t} c_t - v_{l,t} l_t + \beta \omega_{t+1} && (35) \end{aligned}$$

and given $\{\omega_t\}_{t=0}^{\infty}$, k_0. Once again we denote the multiplier on the promise-making constraint (34) by λ_t^m and on the promise-keeping constraint (35) by λ_t^k. The resource multiplier is denoted η_t.

[42]Parameter values are as before: $\beta = 0.96$, $\gamma = 0.024$, $\chi = 0.048$ and $\bar{y} = 0.1$.

[43]Recall that the New Keynesian Phillips Curve (6) is not vertical in the long run, so permanently positive inflation is consistent with output being permanently above its flexible price level.

Necessary optimality conditions with respect to c_t, l_t and k_{t+1} are:

$$
\begin{aligned}
u_{c,t} - \eta_t - \lambda_t^m u_{cc,t} k_{t+1} + \lambda_t^k \left(u_{c,t} + u_{cc,t} c_t \right) &= 0 && (36) \\
-v_{l,t} + \eta_t F_{l,t} - \lambda_t^k \left(v_{l,t} + v_{ll,t} l_t \right) &= 0 && (37) \\
-\eta_t + \beta \eta_{t+1} \left(1 + F_{k,t+1} - \delta \right) - \lambda_t^m u_{c,t} &= 0 && (38)
\end{aligned}
$$

Once more we need a way to fix the promise sequence, and hence place a restriction on λ_t^m and λ_t^k in each time period. Again, to be consistent with RPE policy this must ensure $\beta \left(\lambda_t^m + \lambda_t^k \right)$ approaches λ_t^k in steady state. In this case the choice of restriction is less immediate. Capital accumulation clearly means it is not possible to choose constant values for all endogenous variables in all periods. Instead, we make use of the additively separable structure of the problem over time to arrive at simple, intuitive policy rules.

Equation (37), the first-order condition with respect to l_t, is a cross-restriction linking t-dated variables alone. To the extent that dynamics enter it, these operate only indirectly through the resource multiplier, η_t. Once the capital accumulation decision has determined the shadow value of resources in t, the optimal labour supply decision has effectively been 'orthogonalised' from any dynamic aspects of choice. This is not true of the choice of c_t for arbitrary promise sequences. Condition (36) contains a term in k_{t+1}, which derives from the fact that a change in period-t consumption will affect the period-t value of any capital savings carried from t to $t+1$, and hence the ease of satisfying the promise-keeping constraint. Yet it is possible to place a restriction on the promise-keeping and promise-making multipliers in such a way that the consumption choice is likewise orthogonalised from all variables that influence outcomes at $t+1$. Consider the condition:

$$
\beta \left(\lambda_t^m + \lambda_t^k \right) \omega_{t+1} = \lambda_t^k \omega_t \tag{39}
$$

In words, the t-dated shadow value of the promise made for $t+1$ should equal the t-dated shadow cost of the promise kept in t. This is consistent with our required restriction for steady state, since any steady state must involve the promises in t and $t+1$ converging to one another, leaving a steady-state restriction equivalent to (25).

Condition (39) simplifies considerably when the promise values are replaced by the objects from the inner problem that correspond to them in (34) and (35).[44] In that case, we have:

$$
\lambda_t^m u_{c,t} k_{t+1} = \lambda_t^k \left(u_{c,t} c_t - v_{l,t} l_t \right) \tag{40}
$$

Substituting this into (36) gives:

$$
u_{c,t} - \eta_t + \lambda_t^k \left(u_{c,t} + u_{cc,t} \frac{v_{l,t}}{u_{c,t}} l_t \right) = 0 \tag{41}
$$

We now have an expression that depends on t-dated objects alone. It can be combined with the

[44] That is, setting $\omega_t - \beta\omega_{t+1} = u_{c,t} c_t - v_{l,t} l_t$ and $\beta\omega_t = u_{c,t} k_{t+1}$. Note that using these as equalities is without loss of generality: in the event that constraint (34) and/or (35) is slack, the corresponding multiplier will be zero and the substitution will be of no consequence.

optimality condition on labour supply to eliminate the promise-keeping multiplier λ_t^k, and arrive at an interpretable, time-invariant condition for RPE policy:

$$\frac{u_{c,t} - \eta_t}{u_{c,t} + c_t u_{cc,t} \frac{v_{l,t} l_t}{u_{c,t} c_t}} = \frac{v_{l,t} - \eta_t F_{l,t}}{v_{l,t} + v_{ll,t} l_t} \tag{42}$$

The objects in the numerators here are immediately recognisable as policy 'wedges': on the left-hand side, the gap between the marginal value of consumption resources for the representative agent and the shadow cost of resources; on the right-hand side, the gap between the marginal utility loss from an extra unit of work and the shadow value of the resources it generates.

The objects in the denominators are analogous to marginal revenue terms in textbook monopoly analysis, which in turn depend on inverse elasticities. In the decentralised economy with taxes, the effective post-tax price of consumption in period t is $u_{c,t}$, and the post-tax wage in period t is $v_{l,t}$.[45] An increase in l_t by one unit will increase total wage income by an amount $w_t + l_t \frac{dw_t}{dl_t}$, where w_t denotes the wage rate. This is the denominator on the right-hand side. An increase in c_t has more complicated effects. It increases the period-t cost of consumption by an amount $p_t + c_t \frac{dp_t}{dc_t}$, where p_t denotes the price of consumption, but it will also revalue any capital income that the consumer is holding in period t. A general assessment of the value of this will be time-contingent: revaluing capital income may be desirable after it has been accumulated, but promising to do so beforehand may not. This is the time-inconsistency problem once more. But precisely through imposing restriction (39), we have restricted attention in each period to the revaluation only of that share of consumption that is earned through work in period t: $\frac{v_{l,t} l_t}{u_{c,t} c_t}$, or $\frac{w_t l_t}{p_t c_t}$. This accounts for the final fraction in the denominator on the left-hand side.

Taken together, the condition states that it should not be possible for the policymaker to benefit from a joint change in period-t consumption and labour supply that respects budget balance, holding constant the net value of expenditure from capital income in each period. The left-hand side is the net benefit from increased consumption per unit change in the consumer's budget; the right-hand side is the net cost from increased labour supply.

Finally, it is possible to use conditions (37) and (41) in (38) to obtain a restriction linking the capital, consumption and labour supply wedges in all time periods. We have:

$$k_{t+1} \left[\beta \eta_{t+1} \left(1 + F_{k,t+1} - \delta \right) - \eta_t \right] = c_t \frac{\eta_t - u_{c,t}}{1 - \frac{v_{l,t} l_t}{u_{c,t} c_t} \frac{1}{\varepsilon_t^c}} + l_t \frac{v_{l,t} - \eta_t F_{l,t}}{1 + \frac{1}{\varepsilon_t^l}} \tag{43}$$

where we use ε_t^c to denote the Frisch elasticity of consumption in period t with respect to its price, and ε_t^l to denote the Frisch elasticity of labour supply in t. This equation captures a very intuitive trade-off. The policymaker would like, if possible, to ensure that savings are at a sufficient level to equate the shadow value of capital saved for $t+1$ with its shadow cost at t. The benefit of greater savings is captured by the left-hand side of the expression. But in order to encourage greater savings, consumers must be endowed with sufficiently high disposable income. The higher is disposable income, the higher will be current consumption, and the lower will be labour supply. The net cost of

[45] Clearly this is true up to a normalisation of the price index.

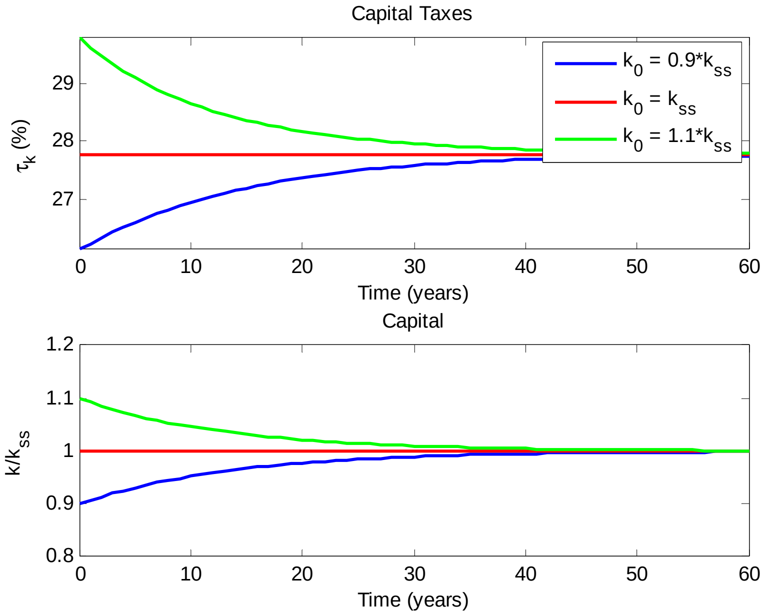

Figure 5: Capital tax rates and capital transition under RPE policy

these two effects is captured, in turn, by the two terms on the right-hand side of the expression. Overall, the policymaker is thus trading off the costs and benefits of providing consumers with spending power. This is the underlying policy problem once time-inconsistent dynamics have been filtered out.

Conditions (42) and (43) together with the resources and implementability condition are jointly sufficient to solve for a dynamic allocation, together with the implied values for policy instruments. Figure 5 charts the evolution of the capital tax rate and the capital stock under differing values for k_0, given this solution approach.[46] By contrast with the extreme time-sensitivity of Ramsey policy, it is evident that RPE tax policy departs from steady state only to the extent that the underlying capital stock does likewise. When capital is above steady state, the policymaker has a reduced incentive to provide wealth to the representative consumer, and capital taxes are relatively high. The converse is true when capital is below steady state: in this case is a greater incentive to encourage accumulation, and capital taxes are low.

Example 3: Insurance with limited commitment

In the third example, the inner problem can be written as:

$$\max_{\left\{c_t(\sigma), c_t^p\right\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left[(1-\mu) \sum_{\sigma=0}^{\infty} (1-p)^{\sigma} pu\left(c_t\left(\sigma\right)\right) + \mu u(c_t^p) \right]$$

[46]Parameter values are as detailed in Section 2.3.

subject to:

$$
\begin{aligned}
(1-\mu)\sum_{\sigma=0}^{\infty}(1-p)^{\sigma}pc_s(\sigma)+\mu c_s^p &\leq [1-(1-\mu)p]y^l+(1-\mu)py^h && (44)\\
V(\sigma) &\leq u(c_t(\sigma))+\beta\mathbb{E}_t\omega_{t+1}(\sigma') && (45)\\
\omega_t(\sigma) &\leq u(c_t(\sigma))+\beta\mathbb{E}_t\omega_{t+1}(\sigma') && (46)
\end{aligned}
$$

and given promise sequences $\{\omega_t(\sigma)\}_{t=0}^{\infty}$. Condition (45) here is the promise-making constraint, and condition (46) is the promise-keeping constraint. Placing multiplier η_t on the resource constraint, $\lambda_t^m(\sigma)$ on the promise-making constraint and $\lambda_t^k(\sigma)$ on the promise-keeping constraint,[47] necessary optimality conditions for this problem are:

$$
\begin{aligned}
u_{c,t}(\sigma)\left(1+\lambda_t^m(\sigma)+\lambda_t^k(\sigma)\right)-\eta_t &= 0 && (47)\\
u_{c,t}^p-\eta_t &= 0 && (48)
\end{aligned}
$$

where $u_{c,t}^p$ is used to denote the marginal utility of consumption in period t for permanently poor agents. Thus λ_t^m and λ_t^k play the usual role in this class of problem, augmenting the policymaker's underlying cross-sectional Pareto weights. Agents whose promise-making or promise-keeping constraints bind in period t will receive a more generous allocation than others.

To progress further, we need to impose a restriction on the relative values of λ_t^m and λ_t^k. Consistent with the results of Proposition 9, in any steady state these must satisfy:

$$
\lambda_{ss}^k(\sigma)=\beta\left[\lambda_{ss}^m(\sigma-1)+\lambda_{ss}^k(\sigma-1)\right] \tag{49}
$$

for $\sigma>0$, with $\lambda_t^m(0)$ set to a sufficiently high value to ensure continued participation by those who have received a positive wealth shock. Again, we are interested in policies that will not exhibit transition dynamics *per se*. In a manner similar to the inflation bias example, this can be done by requiring (49) to hold in every time period. That is, we impose the restriction:

$$
\lambda_t^k(\sigma)=\beta\left[\lambda_t^m(\sigma-1)+\lambda_t^k(\sigma-1)\right] \tag{50}
$$

for all t. This gives us a within-period multiplier recursion, which can be solved backwards to eliminate the terms in λ^k:

$$
\lambda_t^k(\sigma)=\sum_{s=1}^{\sigma}\beta^s\lambda_t^m(\sigma-s) \tag{51}
$$

If we conjecture that the solution will set $\omega_t(\sigma)>V(\sigma)$ for all $\sigma>0$, which is necessary if low types are always to be benefitting from insurance, then $\lambda_t^m(\sigma)=0$ for all $\sigma>0$, and we have:

$$
\lambda_t^k(\sigma)=\beta^{\sigma}\lambda_t^m(0) \tag{52}
$$

Substituting this into (47), and dropping the redundant dependence of λ_t^m on σ, cross-sectional allo-

[47] These latter two multipliers are normalised by relevant population weights – i.e., a factor $(1-\mu)(1-p)^{\sigma}p$.

cations now satisfy:

$$u_{c,t}\left(\sigma\right)=\frac{\eta_t}{1+\beta^{\sigma}\lambda_t^m} \tag{53}$$

This is a time-invariant condition, which will combine with the static resource constraint to deliver a time-invariant consumption distribution and constant resource multiplier in all periods. Under this solution, all agents' consumption allocations are determined by reference to those whose current participation constraint binds. It is instructive to contrast condition (53) with the equivalent term that emerges under Ramsey policy. In that case the following condition is well known to characterise allocations across agents:[48]

$$u_{c,t}\left(\sigma\right)=\frac{\eta_t}{1+\lambda_{t-\sigma}^m} \tag{54}$$

That is, an agent whose participation constraint was last binding in period $t-\sigma$ retains a cross-sectional Pareto weight in t that is equal to the shadow value of promise-making in $t-\sigma$. Augmentations to the cross-sectional weights do not exhibit decay. Yet the total quantity of resources available each period is unchanging. This implies that Pareto weights will be continually bid up over time. A value of λ_t^m that would be sufficient to divert a substantial share of resources to high-income agents when all other Pareto weights are set to 1 will be inadequate to do so when a large share of the population has already experienced a high-income shock in the past. This process causes the resource multiplier to increase over time, to the point where the consumption of permanently low-income types is eventually driven to its autarky lower bound – as set out in section 2.3. In the long run, the policymaker is completely bound by past commitments.

Condition (53) departs from this in two regards. First, instead of placing an explicit dynamic link on promise multipliers from one period to another, it links the cross-sectional distribution of these multipliers, across agents at t with different earnings histories. Second, and more substantively for policy, it allows Pareto weights to decay. This means that past promise commitments never come to dominate allocations in the same manner as under Ramsey policy, and redistribution to permanenetly low-income agents will take place in perpetuity. Figure 6 illustrates by contrasting the consumption of permanently low-income types under RPE and Ramsey policy, based on the same log-utility calibration detailed in Section 2.3. The absence of any downward drift under RPE policy reflects a constant value for the resource multiplier.

Consistent with its formal definition, RPE policy continues to deliver desirable redistribution as time progresses. After a relatively short horizon it comes to deliver higher welfare than optimal choices given the continuation sequence of Ramsey promises. Figure 7 illustrates this. At each point in time it charts the consumption level that would deliver an equivalent value for social welfare to the Ramsey and RPE continuations if this consumption level were allocated uniformly to all agents in the population, in perpetuity. This is normalised relative to the first-best consumption level.[49] The figure indicates that the relative superiority of the Ramsey promise path is reversed after 18 years. After this time, all policymakers would strictly prefer to switch to RPE policy: continuation with Ramsey no longer satisfies the Pareto criterion.

[48] See, for instance, Marcet and Marimon (2015)

[49] That is, relative to the optimal allocation when participation constraints are ignored. A utilitarian policymaker would always share resources euqally in such circumstances.

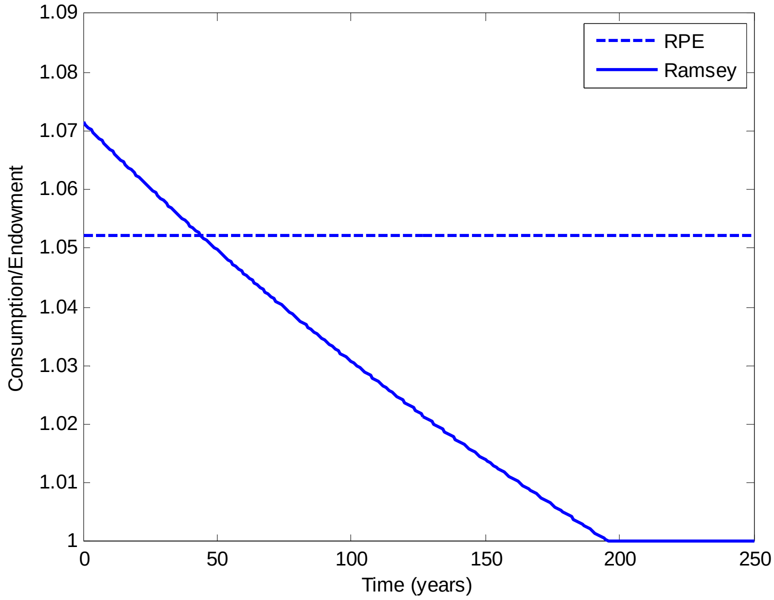

Figure 6: Ramsey and RPE policy: consumption of low-income agents in the limited commitment example

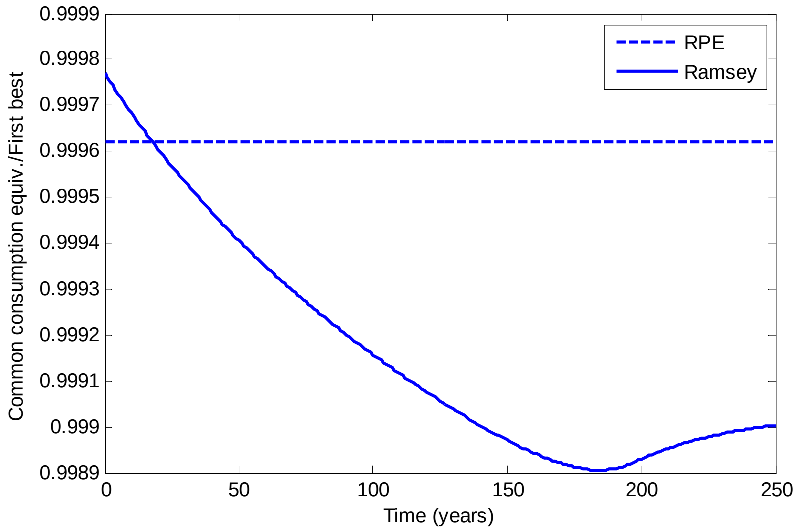

Figure 7: Ramsey and RPE policy: welfare comparisons in the limited commitment example

The RPE allocation has interesting implications for the appropriate structure of policy interventions. Influential work by Alvarez and Jermann (2000) has illustrated how the Ramsey allocation in limited commitment models can be decentralised simply by adding borrowing constraints to a standard Walrasian setting. These constraints prevent agents from realising the value of their earnings in high-income states before these states have been realised. This is a necessary restriction given the scope to default on past debts and consume in autarky. The fact that all other aspects of the Ramsey solution can be decentralised by Arrow-Debreu trades reflects the fact that relative Pareto weights are constant from t to $t+1$ in the Ramsey allocation across all pairs of all agents who do not receive high-income shocks. This means there is no role for savings taxes. Specifically, we can define the *ex-post* marginal savings wedge, τ_{t+1}^{s}, by:

$$\frac{\beta u_{c,t+1}\left(1-\tau_{t+1}^{s}\right)}{u_{c,t}} := \frac{\beta \eta_{t+1}}{\eta_{t}} \tag{55}$$

That is, the tax wedge that is implicitly applied relative to the shadow cost of transferring public funds through time, under any given allocation. Under the Ramsey allocation this wedge is zero whenever the agent does not receive a high-income draw in $t+1$: this follows immediately from (54). Under the RPE allocation it satisfies:

$$\left(1-\tau_{t+1}^{s}\right) = \frac{1+\beta^{\sigma+1}\lambda_{t+1}^{m}}{1+\beta^{\sigma}\lambda_{t}^{m}} \tag{56}$$

In the stationary solution λ_{t}^{m} is constant over time, implying a positive tax wedge that is monotonically decreasing in σ, approaching zero for those whose last high-income shock was many periods in the past. Hence there is effectively a progressive tax on savings. This encourages agents with high-income draws to frontload their consumption, and thereby ensures that prior wealth shocks do not come to dominate the long-run distribution of consumption in the economy.

7. RPE policy and promise trading `[in progress]`

There is a very close relationship between the RPE policies that we have introduced and the market interactions among successive generations that have been studied in detail in overlapping generations models from Samuelson (1958) onwards. The purpose of this brief section is to highlight these similarities, since the parallels can help to clarify our earlier analysis.

Consider a textbook two-period, pure-endowment OLG economy in which agents' endowments are higher when they are young than when they are old. These endowments are constant over time, and denoted y^{y} and y^{o}. Agents face an infinite array off intertemporal consumption prices, denoted $\{p_{t}\}_{t=0}^{\infty}$, at which they are free to trade their endowments. These trades are conducted to maximise their *ex-ante* welfare, assessed according to a standard utility function:

$$U_{t} := u\left(c_{t}^{y}\right) + u\left(c_{t+1}^{o}\right)$$

The budget constraint is thus:

$$p_t \left(c_t^y - y^y\right) + p_{t+1} \left(c_{t+1}^o - y^o\right) \leq 0 \tag{57}$$

Market clearing requires that the outcome of this choice problem for all $t \geq 0$ should satisfy $c_t^y + c_t^o = y^y + y^o$, with c_0^o the consumption level for the initial old generation, which is unrestricted.

This problem can be recast as a choice of promises over time, where the promises in question are intergenerational transfers. Denoting by ω_t^k the transfer that is delivered through market trades by a generation when it is young, and ω_{t+1}^m the transfer that the same generation anticipates when it is old, we have:

$$\omega_t^k := y^y - c_t^y \tag{58}$$

$$\omega_{t+1}^m := c_{t+1}^o - y^o \tag{59}$$

Trivially, the problem for each generation can be mapped into one of choosing promises to maximise a value function expressed over these promises, subject to a linear price vector:

$$\max_{\left\{\omega_t^k, \omega_{t+1}^m\right\}} V\left(\omega_t^k, \omega_{t+1}^m\right) := u\left(y^y - \omega_t^k\right) + u\left(y^o + \omega_{t+1}^m\right)$$

subject to:

$$p_{t+1}\omega_{t+1}^m \leq p_t\omega_t^k \tag{60}$$

Market clearing then requires that $\omega_t^k = \omega_t^m$ for all $t > 0$. The transfer to the initial generation of old agents, ω_0^k, is not restricted by a market clearing condition. The efficiency properties of the V function in the choice of promise vectors are very similar to the properties of the value function for the inner problem that we defined in Section 3, $V\left(\{\omega_t\}_{t=0}^{\infty}, x_0\right)$.

A Pareto efficient sequence of promises can be defined in the same manner as in Section 4 above, with the same distinction between weak and strict definitions. Here too it will not be possible to satisfy the strict definition recursively, since each new young generation will always benefit from eliminating any transfers to the initial old.[50] A weakly Pareto efficient promise sequence from period τ onwards is then an array $\{\omega_t\}_{t=\tau}^{\infty}$ such that there does not exist an alternative sequence $\{\omega_t'\}_{\tau}^{\infty}$ with the property:

$$V\left(\omega_t', \omega_{t+1}'\right) - V\left(\omega_t, \omega_{t+1}\right) \geq \varepsilon > 0$$

for all $t \geq \tau$. A sequence is recursively Pareto efficient if (and only if) it satisfies this definition for all $\tau \geq 0$.

In this setting, a promise sequence that induces convergence in allocations to a steady state can be consistent with recursive Pareto efficiency only if marginal utilities satisfy the following relationship in steady state:

$$u'\left(c_{ss}^y\right) = u'\left(c_{ss}^o\right) \tag{61}$$

If convergence were to any other steady state, a uniform increase or reduction in promises relative

[50]This reflects the well-known fact that the core is empty in OLG models.

to their steady-state values could strictly increase the welfare of all generations, once allocations had come sufficiently close to steady state. As is well known, this efficient allocation is decentralised by a price sequence with the property that the real interest rate converges to one, i.e.:

$$\lim_{t\to\infty}\left(\frac{p_{t+1}}{p_t}\right)=1 \tag{62}$$

Convergence to this steady state is not the only possible outcome that could satisfy recursive Pareto efficiency. In precisely this setting, Grandmont (1985) highlighted the possibility of Walrasian allocations that exhibit limit cycles in the real interest rate, and thus in allocations, when income effects are large enough. These cycles cannot be Pareto-dominated for all generations at the limit. What *is* required for recursive Pareto efficiency is that there should never come a time period τ such that either the inequality:

$$u'\left(c_t^y\right)-u'\left(c_t^o\right)\geq\varepsilon>0$$

holds for all $t\geq\tau$, or the converse:

$$u'\left(c_t^y\right)-u'\left(c_t^o\right)\leq-\varepsilon<0$$

for all $t\geq\tau$. This is a direct analogue of Proposition 9 above.

The similarities between RPE outcomes in the OLG and Kydland-Prescott already appear substantial. In both cases it is only a weak definition of Pareto efficiency that can be applied recursively, and in both cases the RPE requirement places a restriction on long-run outcomes alone, not the transition.

We can go further by examining an analogue to the OLG market mechanism as a 'decentralisation' scheme for RPE allocations in Kydland and Prescott problems. The promise objects featuring in constraints (17) and (18) for any given t can be divided into promises made at t for $t+1$, ω_{t+1}^m, and promises kept at t, ω_t^k.[51] There is no technical barrier to then studying a variant of the inner problem in which separate sequences $\left\{\omega_{t+1}^m\right\}_{t=0}^{\infty}$ and $\left\{\omega_t^k\right\}_{t=0}^{\infty}$ are applied to promise-making and promise-keeping sides of the restrictions respectively. The value of this problem in period 0 can be denoted $V\left(\left\{\omega_t^k,\omega_{t+1}^m\right\}_{t=0}^{\infty},x_0\right)$. If $\omega_t^k=\omega_t^m$ for all $t>0$ then this object is equal to the value function that we have already used extensively, $V\left(\left\{\omega_t\right\}_{t=0}^{\infty},x_0\right)$. When the constraints are of the infinite-horizon form, the shadow cost from increasing ω_t^k by a unit is $\beta^t\lambda_t^k$, and the shadow benefit from increasing ω_{t+1}^m by a unit is $\beta^{t+1}\left(\lambda_t^k+\lambda_t^m\right)$, where λ_t^k and λ_t^m are the multipliers on the period-t promise-keeping and promise-making constraints respectively.

Consider now a sequential market trading mechanism, whereby the policymaker in each period t has the capacity to choose ω_t^k and ω_{t+1}^m, subject to a linear budget constraint:

$$p_{t+1}\left(\omega_{t+1}^m-\bar{\omega}_{t+1}\right)\leq p_t\left(\omega_t^k-\bar{\omega}_t\right) \tag{63}$$

The objects $\left\{\bar{\omega}_t\right\}_{t=0}^{\infty}$ can be treated as 'endowment' values for the promises: values that would obtain in the absence of any trade. In the OLG setting, where the promises correspond to intergenerational transfers, the equivalent endowment values are zero – hence their absence from the equivalent con-

[51] The ω_{t+1}^m objects enter on the left-hand side of (17) and (18), and the ω_t^k objects on the right-hand side of (18) alone.

straint (60).

Given an endowment sequence $\{\bar{\omega}_t\}_{t=0}^{\infty}$ and an initial state vector x_0, we can define a promise-trading equilibrium as a sequence of promises $\left\{\omega_t^{k*}, \omega_{t+1}^{m*}\right\}_{t=0}^{\infty}$, prices $\{p_t\}_{t=0}^{\infty}$ and endogenous states $\left\{x_{t+1}^{*}\right\}_{t=0}^{\infty}$ such that:

1. There is market clearing in promises: $\omega_t^{k*} = \omega_t^{m*}$ for all $t > 0$

2. Choices are sequentially optimal for each policymaker:

$$\left\{\omega_t^{k*}, \omega_{t+1}^{m*}\right\} \in \arg \max_{\left\{\omega_t^{k}, \omega_{t+1}^{m}\right\}} V\left(\left\{\omega_s^{k}, \omega_{s+1}^{m}\right\}_{s=t}^{\infty}, x_t^{*}\right)$$

   subject to (63) and given $\omega_s^k = \omega_s^{k*}$ and $\omega_{s+1}^m = \omega_{s+1}^{m*}$ for all $s > t$.

3. The endogenous states $\left\{x_{t+1}^{*}\right\}_{t=0}^{\infty}$ solve the inner problem given $\left\{\omega_t^{k*}, \omega_{t+1}^{m*}\right\}_{t=0}^{\infty}$.

As an equilibrium concept this is almost identical to the Walrasian mechanism analysed in the OLG setting above, the only major complication being the need to allow for endogenous evolution of the state vector.

Optimal choice by the policymaker in each time period implies the following condition:

$$\frac{p_{t+1}}{p_t} = \frac{\beta\left(\lambda_t^k + \lambda_t^m\right)}{\lambda_t^k} \tag{64}$$

From Proposition 9 and its Corollary, we know that if an allocation is recursively Pareto efficient then the fraction on the right-hand side of (64) cannot converge to any value other than unity. This implies, exactly as in the OLG case, that the intertemporal relative price of promises must converge to one, if it converges at all.

The analogue to the OLG setting sheds important light on the question of transition dynamics that were studied in our three specific examples in Section 6. There, we noted that the RPE desirability concept alone was not sufficient to deliver a unique transition, though in each case it was possible to find a restriction on the multipliers that ensured the final policy outcome took a simple, time-invariant form. Resolving the multiplicity is equivalent to fixing a particular sequence for the 'endowment' promises $\{\bar{\omega}_t\}_{t=0}^{\infty}$. To see this, note that (63) and (64) will combine to give:

$$\beta\left(\lambda_t^k + \lambda_t^m\right)\left(\omega_{t+1}^m - \bar{\omega}_{t+1}\right) = \lambda_t^k\left(\omega_t^k - \bar{\omega}_t\right) \tag{65}$$

where equality follows from complementary slackness.

It is precisely a restriction of this form that was necessary to obtain simple policy rules in our three examples.[52] In the inflation bias and limited commitment models, we implicitly selected a value for $\bar{\omega}_t$ in each period that was consistent with an equilibrium supporting price ratio, $\frac{p_{t+1}}{p_t}$, equal to one. In the capital tax example, we selected a value for $\bar{\omega}_t$ equal to zero. Both cases could be seen as

[52] Whether these rules will in turn be consistent with a *unique* RPE outcome is an additional question that we leave unanswered. As noted above, Pareto efficient limit cycles are possible in the OLG setting, so there seems no good reason to rule them out in general Kydland and Prescott problems.

ways to give equal treatment to different periods' policymakers – either by facing them with common prices, or providing them with common endowments. Which of these provides the more appealing normative solution may well vary from case to case. It may additionally be possible to provide a positive motivation for selecting certain arrays of endowments, for instance based on the promises that would implicitly obtain in an inefficient Markovian equilibrium. This seems an interesting area for further work.

8. 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 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. 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 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. The transition dynamics of policy depend only on the economy's natural state vector, and long-run outcomes are 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 that of the private sector. 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. The relationship between policy that is recursively Pareto efficient and policy that is sustainable in a reputational equilibrium is a priority for future work.

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. This is left for future 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] Chang, R. (1998), 'Credible Monetary Policy with Long-Lived Agents: Recursive Approaches', *Journal of Economic Theory*, 81, 431–461.

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

[16] 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.

[17] 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.

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

[19] 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.

[20] 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.

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

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

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

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

[25] Grandmont, J-M. (1985), 'On Endogenous Competitive Business Cycles', *Econometrica*, 53(5), 995–1045.

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

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

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

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

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

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

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

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

[34] 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.

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

[36] Marcet, A., and R. Marimon (2015), 'Recursive Contracts', manuscript.

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

[38] 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.

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

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

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

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

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

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

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

[46] Straub, L., and I. Werning (2015), 'Positive Long Run Capital Taxation: Chamley-Judd Revisited', manuscript.

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

[48] 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.

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

[50] 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.

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

[52] 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

The spaces $\Omega(x_t)$ and $(A \times X)^\infty$ contain bounded sequences, and thus are l^∞ metric spaces when endowed with the sup norm. Quasi-concavity of h and p, quasi-convexity of h_0, linearity of l, and the linear presence of the promise values in (17) and (18) implies that the constraint set for the inner problem is convex-valued and continuous in $\{x_t, \{\omega_s\}_{s=t}^\infty\}$ at all interior points of $\Omega(x_t)$. The objective function r is continuous under Assumption 1 and quasi-concave under Assumption 5. The result then follows by a direct application of Berge's Theorem of the Maximum.

A.2. Proof of Proposition 2

Suppose that the statement were not true. Then for some $\tau \geq 1$ there must exist a continuation allocation $\{x''_{s+1}, a''_s\}_{s=t+\tau}^\infty$ that satisfies the constraints of the inner problem from τ onwards and delivers a higher value $W_{t+\tau}$ than the allocation $\{x'_{s+1}, a'_s\}_{s=t+\tau}^\infty$. But if this is true then the combined allocation $\left\{\{x'_{s+1}, a'_s\}_{s=t}^{t+\tau-1}, \{x''_{s+1}, a''_s\}_{s=t+\tau}^\infty\right\}$ must in turn deliver a higher value for W_t than $\{x'_{s+1}, a'_s\}_{s=t}^\infty$. By the optimality of $\{x'_{s+1}, a'_s\}_{s=t}^\infty$ this cannot be possible unless the combined allocation violates one of the constraints: (2), (3), or (17) and (18) in some period $s \geq t$. The component $\{x'_{s+1}, a'_s\}_{s=t}^{t+\tau-1}$ 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 $\{x'_{s+1}, a'_s\}_{s=t}^\infty$, 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.3. Proof of Proposition 3

The arguments are similar to the proof of Proposition 1. Quasi-concavity of h and quasi-convexity of h_0, together with the linear presence of the promise values in (17) and (18) implies that the constraint set for the inner problem is continuous in $\{\omega_s\}_{s=t}^{\infty}$ at all interior points of $\Omega(x_t)$. Linearity of l implies an equivalent continuity in x_t. The objective function r is continuous under Assumption 1, so the results again follow by a standard application of Berge's Theorem of the Maximum.

A.4. LICQ and Proof of Proposition 4

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 $2j$ constraints in (17) and (18) 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 (17) and (18). 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 (vector-valued) h_0 and h functions index only the subset of constraints of types (17) and (18) that bind.

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 4

Strict concavity in the return function r and the convexity of the constraint set implied by Assumptions 3 and 5 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 $\{\omega_s\}_{s=t}^{\infty}$.[53] Associated with this solution is a standard set of Kuhn-Tucker conditions. Within period s and for stochastic

[53]C.f. the proof of Proposition 3.

draw σ_s, these conditions with respect to a_s will take the form:

$$\begin{aligned} \frac{\partial r}{\partial a_s} + \lambda_s'(\sigma_s)\frac{\partial G(a_s,\sigma_s)}{\partial a_s} &= 0 &(66)\\ \lambda_s'(\sigma_s)\,G(a_s,\sigma_s) &= 0 &(67)\\ \lambda_s'(\sigma_s) &\geq 0 &(68)\end{aligned}$$

where $G(a_s,\sigma_s)$ stacks the constraints in (17) and (18):

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

and $\lambda_s = \left[(\lambda_s^m)', (\lambda_s^k)'\right]'$ is the stacked vector of multipliers on these constraints.[54] 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.5. Proof of Proposition 5

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:

$$V\left(\{\alpha\omega_s' + (1-\alpha)\,\omega_s''\}_{s=t}^{\infty}, x_t\right) \geq \bar{V} \tag{69}$$

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

Let $\mathbf{y}' := \{x_{s+1}', a_s'\}_{s=t}^{\infty}$ and $\mathbf{y}'' := \{x_{s+1}'', a_s''\}_{s=t}^{\infty}$ 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 5) that (69) 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, (69) 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 (17) and (18) are also satisfied. Consider (17). We need:

$$\begin{aligned} h\left(\alpha a_s' + (1-\alpha)\,a_s''\right) + \beta\left[\alpha\omega_{s+1} + (1-\alpha)\,\omega_{s+1}\right] &\geq \\ h_0\left(\alpha a_s' + (1-\alpha)\,a_s'', \alpha x_s' + (1-\alpha)\,x_s'', \alpha x_{s+1}' + (1-\alpha)\,x_{s+1}''\right) \end{aligned}$$

[54] To ease on notation we assume in writing (66) 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 (66).

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

$$\begin{aligned}\alpha h\left(a_s'\right)+(1-\alpha)\, h\left(a_s''\right)+\beta\left[\alpha\omega_{s+1}+(1-\alpha)\,\omega_{s+1}\right] &\geq \\ \alpha h_0\left(a_s',x_s',x_{s+1}'\right)+(1-\alpha)\, h_0\left(a_s',x_s',x_{s+1}'\right)\end{aligned}$$

But by concavity of h:

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

and by convexity of h_0:

$$\begin{aligned}h_0\left(\alpha a_s'+(1-\alpha)\, a_s'',\alpha x_s'+(1-\alpha)\, x_s'',\alpha x_{s+1}'+(1-\alpha)\, x_{s+1}''\right) &\leq \\ \alpha h_0\left(a_s',x_s',x_{s+1}'\right)+(1-\alpha)\, h_0\left(a_s',x_s',x_{s+1}'\right)\end{aligned}$$

establishing the desired inequality. An identical argument confirms that (18) is 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 $\left\{\alpha\omega_s'+(1-\alpha)\,\omega_s''\right\}_{s=t}^{\infty}$, completing the proof.

A.6. Proof of Proposition 7

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. Let ω_τ be a promise vector that constrains the inner problem, and denote by λ_s^m and λ_s^k the vectors of current-value multipliers on constraints (17) and (18) respectively for generic period s.[55] These multipliers are strictly positive whenever the respective constraints bind. The marginal effect of increasing ω_τ on $V\left(\left\{\omega_s^R\right\}_{s=t}^{\infty},x_t\right)$ is $\beta^{\tau-t}\left(\lambda_{\tau-1}^m+\lambda_{\tau-1}^k-\lambda_\tau^k\right)$ for $\tau>t$, and $-\lambda_\tau^k$ for $\tau=t$, as shown in Proposition 4. Suppose first that the only binding constraint is (17), 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.[56] Suppose instead that constraint (18) 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.7. Proof of Proposition 8

A.7.1. Preliminaries

Let $\bar{\omega}>0$ be a uniform bound on $\left\{\omega_s\right\}_s$ (this exists by the Assumption in section 3.1.1), i.e. $\|\omega_s\|<\bar{\omega}$ for all s and all $\left\{\omega_s\right\}_s$ sequences in Ω, where $\|\cdot\|$ is the sup norm. Our analysis will focus on a transformation of the underlying promise sequence, given by $\omega^\beta:=\left\{\beta^s\omega_s\right\}_{s=0}^{\infty}$. If we let ω denote complete promise sequences, clearly we can recover an ω associated with any given ω^β by multiplying through β^{-s} for all s. Denote this transformation $\beta^{-1}\left(\omega^\beta\right)$. That is, $\omega=\beta^{-1}\left(\omega^\beta\right)$. The use of ω^β in place of

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

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

ω is for analytical purposes. In particular, ω^{β} lives in the space:

$$S_{\bar{\omega}} = \left\{ y \in (\mathbb{R}^{j+k} \times \mathbb{R}^{\varsigma})^{\infty} : \|y_s\| \le \beta^s \bar{\omega} \right\}$$

Since $\beta^s \bar{\omega} \to 0$, this space is well-known to be compact. Compactness is an important property in enabling fixed-point theorems to be applied.[57] Recall that the set $\Omega(x_0)$ denotes the promise sequences for which the interior problem's constraint set is non-empty. This is closed for any x_0 . We can denote by $\Omega^{\beta}(x_0)$ the corresponding set of ω^{β}:

$$\Omega^{\beta}(x_0) := \left\{ \omega^{\beta} \in S_{\bar{\omega}} : \beta^{-1}\left(\omega^{\beta}\right) \in \Omega(x_0) \right\}$$

Since $\Omega(x_0)$ is closed, so too is $\Omega^{\beta}(x_0)$. Hence $\Omega^{\beta}(x_0)$ is a closed subset of the compact set $S_{\bar{\omega}}$. It follows that $\Omega^{\beta}(x_0)$ is also compact. By assumption $\Omega(x_0)$ is convex, and $\Omega^{\beta}(x_0)$ inherits this property.

A.7.2. Construction of a continuous mapping

The proof works by constructing a continuous mapping from $\Omega^{\beta}(x_0)$ into itself, with the property that if ω^{β} is a fixed point of this mapping, then the promise sequence $\omega = \beta^{-1}\left(\omega^{\beta}\right)$ satisfies recursuive weak Pareto efficiency.

Step 1: Define a Pareto-dominating correspondence

Fix some $\omega^{\beta} \in \Omega^{\beta}(x_0)$, with corresponding $\omega = \beta^{-1}\left(\omega^{\beta}\right)$. Let $t(\omega)$ be the first time period in which ω is strictly Pareto-dominated, i.e. the lowest value of t such that:

$$V\left(\{\omega'_s\}_{s=\tau}^{\infty}; x_{\tau}^*\right) - V\left(\{\omega_s\}_{s=\tau}^{\infty}; x_{\tau}^*\right) > \varepsilon > 0 \tag{70}$$

for all $\tau \ge t$, some $\varepsilon > 0$ and some $\{\omega'_s\}_{s=t}^{\infty}$. If ω satisfies recursive Pareto efficiency, let $t(\omega) \equiv \infty$.

For all ω such that $t(\omega) < \infty$, define the correspondence $\mathring{\Gamma} : \Omega^{\beta}(x_0) \to \Omega^{\beta}(x_0)$ as equal to the set of $\tilde{\omega}^{\beta}$ such that $\tilde{\omega}$ strictly Pareto-dominates ω in period $t(\omega)$, as in equation (70). Additionally, in these cases let $\Gamma : \Omega^{\beta}(x_0) \to \Omega^{\beta}(x_0)$ be the correspondence whose graph is the closure of the graph of $\mathring{\Gamma}$. Note that by the continuity of V and of $\{x_s^*\}_{s=0}^{\infty}$ in ω, all elements of $\Gamma\left(\omega^{\beta}\right)$ must Pareto-dominate ω. That is, for all $\tilde{\omega}$ such that the corresponding $\tilde{\omega}^{\beta}$ is in $\Gamma\left(\tilde{\omega}^{\beta}\right)$, it follows from (70) that we must have:

$$V\left(\{\tilde{\omega}_s\}_{s=\tau}^{\infty}; x_{\tau}^*\right) \ge V\left(\{\omega_s\}_{s=\tau}^{\infty}; x_{\tau}^*\right) \tag{71}$$

for all $\tau \ge t(\omega)$. When $t(\omega) = \infty$, we simply let $\Gamma\left(\omega^{\beta}\right) = \omega^{\beta}$.

The correspondence $\Gamma : \Omega^{\beta}(x_0) \to \Omega^{\beta}(x_0)$ constructed in this way satisfies the following properties:

1. Convexity (this follows from the quasi-concavity of V)

[57] If ς is countably infinite a similar 'discounting' transformation can be applied to the components of each ω_s to ensure compactness.

2. Upper hemi-continuity (follows from fact Γ has a closed graph)

3. $\omega^\beta \in \Gamma(\omega^\beta)$ for all $\omega^\beta \in \Omega^\beta(x_0)$ (this follows from strict quasi-concavity of V – take a convex combination between any point in $\Gamma(\omega^\beta)$ and ω^β: all agents must strictly prefer all such combinations to ω^β)

Property 3. here implies $\Gamma(\omega^\beta)$ cannot be used directly as the basis for a fixed-point argument, since $\omega^\beta \in \Gamma(\omega^\beta)$ whether or not ω^β satisfies RPE. The idea is instead to construct a correspondence that lives in $\Gamma(\omega^\beta)$ for all ω^β, but contains ω^β only when RPE *is* satisfied.

Step 2: Define a max correspondence within Γ

Let the correspondence $\mathring{\Psi} : \Omega^\beta(x_0) \to \Omega^\beta(x_0)$ be defined as follows:

$$\mathring{\Psi}\left(\omega^\beta\right) := \arg\max_{\tilde{\omega}^\beta \in \Gamma(\omega^\beta)} \left\|\tilde{\omega}^\beta - \omega^\beta\right\|$$

where $\|\cdot\|$ here is the l^2 norm.[58] For a given ω^β this norm defines a continuous function on the compact set $\Gamma(\omega^\beta)$, so existence of the maximum is assured by Wierstrass's theorem. As with $\mathring{\Gamma}$ and Γ, let $\Psi : \Omega^\beta(x_0) \to \Omega^\beta(x_0)$ be the correspondence whose graph is the closure of the graph of $\mathring{\Psi}$. Since $Graph(\mathring{\Psi}) \subset Graph(\Gamma)$ and $Graph(\Gamma)$ is closed, it follows that the closure points of $Graph(\mathring{\Psi})$ must also be contained within $Graph(\Gamma)$. Hence for all $\tilde{\omega}^\beta \in \Psi(\omega^\beta)$, condition (71) must hold.

Finally, denote by $\Psi^c : \Omega^\beta(x_0) \to \Omega^\beta(x_0)$ the correspondence that maps the convex hull of $\Psi(\omega^\beta)$ for all $\omega^\beta \in \Omega(x_0)$. That is, $\tilde{\omega}^\beta \in \Psi^c(\omega^\beta)$ if and only if there is an $\alpha \in [0,1]$ such that $\tilde{\omega}^\beta = \alpha\omega_1^\beta + (1-\alpha)\omega_2^\beta$ for some ω_1^β and ω_2^β in $\Psi(\omega^\beta)$.

We will argue that a fixed point of Ψ^c must satisfy recursive Pareto efficiency. Ψ^c satisfies the following properties:

1. Convexity (by construction)

2. Upper hemi-continuity (closed graph)

A.7.3. Application of Kakutani's fixed-point theorem to Ψ^c

Ψ^c is an upper-hemicontinuous correspondence mapping the convex, compact space $\Omega(x_0)$ into itself. It is non-empty for all $\omega^\beta \in \Omega(x_0)$ (this follows from the fact $\Gamma(\omega^\beta)$ is non-empty by construction) and convex for all $\omega^\beta \in \Omega(x_0)$. Hence, by Kakutani's theorem, Ψ^c has a fixed point.

A.7.4. Proof that this fixed point must satisfy recursive Pareto efficiency

Consider a fixed point of Ψ^c. We need to show that this point, ω^β, must satisfy recursive Pareto efficiency. Suppose otherwise. Then there must exist a point $\tilde{\omega}^\beta$ bounded away from ω^β such that (70) holds for all τ sufficiently large. Since $\|\tilde{\omega}^\beta - \omega^\beta\| > \delta > 0$, it follows that $\omega^\beta \notin \mathring{\Psi}(\omega^\beta)$. It follows

[58] Square summability follows from the fact ω_s is bounded and of finite dimension for all s and $\beta < 1$. The proof extends directly to cases in which ω_s is infinite-dimension but square-summable.

that either ω^β must live in the boundary of $\mathring{\Psi}\left(\omega^\beta\right)$ (i.e., in $\Psi\left(\omega^\beta\right)$), or ω^β is a convex combination of elements of $\Psi\left(\omega^\beta\right)$. Suppose first that $\omega^\beta \in \Psi\left(\omega^\beta\right)$. It follows that there exist ω_1^β and ω_2^β both within an ε ball of ω^β (for any $\varepsilon > 0$) such that $\omega_2^\beta \in \mathring{\Psi}\left(\omega_1^\beta\right)$. But by the continuity of V and of $\{x_s^*\}_{s=0}^\infty$ in ω, we must also have that for ω_1^β sufficiently close to ω^β, $\tilde{\omega}^\beta \in \mathring{\Gamma}\left(\omega_1^\beta\right)$. From the definition of $\mathring{\Psi}\left(\omega^\beta\right)$ it follows that:

$$\left\|\tilde{\omega}^\beta - \omega_1^\beta\right\| \leq \left\|\omega_2^\beta - \omega_1^\beta\right\|$$

As ω_1^β and ω_2^β become arbitrarily close to ω^β the right-hand side of this expression approaches zero, whereas the left-hand side is bounded above zero. This is a contradiction.

The only remaining possibility is that ω^β is a convex combination of two different elements ω_a^β and ω_b^β that are in $\Psi\left(\omega^\beta\right)$ – i.e. $\omega^\beta \in \Psi^c\left(\omega^\beta\right)$ but $\omega^\beta \notin \Psi\left(\omega^\beta\right)$. As shown in Step 2 above, for all $\tilde{\omega}^\beta \in \Psi\left(\omega^\beta\right)$ condition (71) applies. Hence there exists a sufficiently large t such that both:

$$V\left(\{\omega_{a,s}\}_{s=\tau}^\infty; x_\tau^*\right) \geq V\left(\{\omega_s\}_{s=\tau}^\infty; x_\tau^*\right)$$

and

$$V\left(\{\omega_{b,s}\}_{s=\tau}^\infty; x_\tau^*\right) \geq V\left(\{\omega_s\}_{s=\tau}^\infty; x_\tau^*\right)$$

hold. By the strict quasi-concavity of V in ω, it then follows that either $\{\omega_{a,s}\}_{s=\tau}^\infty = \{\omega_{b,s}\}_{s=\tau}^\infty = \{\omega_s\}_{s=\tau}^\infty$ for all $\tau \geq t$ or:

$$V\left(\{\omega_s\}_{s=t}^\infty; x_t^*\right) > V\left(\{\omega_s\}_{s=t}^\infty; x_t^*\right)$$

The latter is clearly a contradiction. It remains to rule out the former – i.e., that the elements of ω_a^β, ω_b^β and ω^β coincide for t sufficiently large. Suppose this were true. By a maintained assumption (to be contradicted) ω^β is Pareto-dominated by some $\tilde{\omega}^\beta$ after a sufficient amount of time. This implies neither ω_a^β nor ω_b^β can be in $\mathring{\Psi}\left(\omega^\beta\right)$. To see why, consider the composite sequence given by $\left\{\left\{\omega_{a,s}^\beta\right\}_{s=0}^{t-1}, \left\{\tilde{\omega}_s^\beta\right\}_{s=t}^\infty\right\}$. By construction this must also Pareto-dominate ω^β after a sufficient amount of time. Under the l^2 norm it is also further away from ω^β than ω_a^β, since strict Pareto-dominance requires that there should always exist an $s > t$ such that $\left\|\tilde{\omega}_s^\beta - \omega_s^\beta\right\| > 0$ for all $t \geq 0$. Hence, by the definiton of $\mathring{\Psi}$, $\omega_a^\beta \notin \mathring{\Psi}\left(\omega^\beta\right)$.

It follows that ω_a^β and ω_b^β are in the boundary of $\Psi\left(\omega^\beta\right)$. Thus for any $\varepsilon > 0$ there exist ω_ε^β and $\omega_{a\varepsilon}^\beta$ within an ε ball of ω^β and ω_a^β respectively such that $\omega_{a\varepsilon}^\beta \in \mathring{\Psi}\left(\omega_\varepsilon^\beta\right)$. But by continuity of V and $\{x_s^*\}_{s=0}^\infty$ in ω, we must have $\tilde{\omega}^\beta \in \Gamma\left(\omega_\varepsilon^\beta\right)$ for sufficiently small ε. This implies the composite series $\left\{\left\{\omega_{a\varepsilon,s}^\beta\right\}_{s=0}^{t-1}, \left\{\tilde{\omega}_s^\beta\right\}_{s=t}^\infty\right\}$ must also be in $\Gamma\left(\omega_\varepsilon^\beta\right)$. But as $\varepsilon \to 0$ the distance between this composite and ω_ε^β (in the l^2 norm) must likewise exceed $\left\|\omega_{a\varepsilon}^\beta - \omega_\varepsilon^\beta\right\|$. This contradicts $\omega_{a\varepsilon}^\beta \in \mathring{\Psi}\left(\omega_\varepsilon^\beta\right)$.

Hence $\omega^\beta \in \Psi^c\left(\omega^\beta\right)$ is inconsistent with the existence of a point $\tilde{\omega}^\beta$ that strictly Pareto dominates ω^β after a sufficient amount of time. It follows that any fixed point of Ψ^c is an ω^β sequence whose counterpart ω satisfies recursive Pareto efficiency.

A.8. Proof of Proposition 9

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

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

such that the value of the object Δ_s, defined as:[59]

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

is either bounded above zero for all $s \geq \tau$ or bounded below zero for all $s \geq \tau$, for some $\tau \geq 0$.

Consider the sequence that sets $\frac{d\omega_s(\sigma')}{d\theta} = 1$ for some $\sigma' \in \Sigma$ and all $s \geq \tau$, 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 $\Delta_s\left(\sigma'\right)$:

$$\Delta_s\left(\sigma'\right) = P\left(\sigma'\right) \sum_{t=s}^{\infty} \beta^{t-s} \left[\sum_{\sigma \in \Sigma} \frac{P\left(\sigma' | \sigma\right) P\left(\sigma\right)}{P\left(\sigma'\right)} \beta \left[\lambda_t^m \left(\sigma\right) + \lambda_t^k \left(\sigma\right) \right] - \lambda_t^k \left(\sigma'\right) \right]$$

If the first inequality stated in the proposition holds, we will have $\Delta_s \geq \frac{1}{1-\beta}\varepsilon > 0$ for all $s \geq \tau$. If the second inequality holds, we will have $\Delta_s \leq -\frac{1}{1-\beta}\varepsilon$. In each case, a strict Pareto improvement is thus possible for all $s \geq \tau$, and RPE cannot be satisfied.

A.8.1. Proof of Corollary to Proposition 9, 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 $\left\{\omega_s\right\}_{s=t}^{\infty}$ 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 (25) holds, there cannot exist a bounded sequence of marginal changes to the promise vectors:

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

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:

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

[59] $P\left(\sigma_{\tau+1} | \sigma_\tau\right)$ here is the transition probability between σ_τ and $\sigma_{\tau+1}$.

for some $\delta > 0$. We know from (25) that:

$$\beta \sum_{\sigma \in \Sigma} \frac{P(\sigma'|\sigma) P(\sigma)}{P(\sigma')} \left[\lambda_{ss}^{m}(\sigma) + \lambda_{ss}^{k}(\sigma)\right] - \lambda_{ss}^{k}(\sigma') = 0$$

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

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

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.