



Department of Economics Discussion Paper Series

Time-Consistently Undominated Policies

Charles Brendon and Martin Ellison

Number 844
January, 2017

Time-Consistently Undominated Policies*

Charles Brendon[†]

Martin Ellison[‡]

January 17, 2018

Abstract

This paper proposes and characterises a new normative solution concept for Kydland and Prescott problems, allowing for a commitment device. A policy choice is dominated if either (a) an alternative exists that is superior to it in a time-consistent subdomain of the constraint set, or (b) an alternative exists that Pareto-dominates it over time. Policies may be time-consistently undominated where time-consistent *optimality* is not possible. We derive necessary and sufficient conditions for this to be true, and show that these are equivalent to a straightforward but significant change to the first-order conditions that apply under Ramsey policy. Time-consistently undominated policies are an order of magnitude simpler than Ramsey choice, whilst retaining normative appeal. This is illustrated across a range of examples.

Keywords: Time Consistency; Undominated Policy; Ramsey Policy

JEL Codes: D02, E61

*This is a new paper, superseding work that previously circulated under the titles ‘Optimal Policy behind a Veil of Ignorance’ and ‘Time-Consistent Institutional Design’. We thank Árpád Ábrahám, Paul Beaudry, Vasco Carvalho, Giancarlo Corsetti, Tatiana Damjanovic, Mike Golosov, Paul Levine, Albert Marcet, Ramon Marimon, Alex Mennuni, Meg Meyer, Michel Normandin, Ctirad Slavik, Rick Van der Ploeg and Simon Wren-Lewis for detailed discussions and comments, together with numerous seminar participants. Brendon is grateful for support from the ADEMU project, ‘A Dynamic Economic and Monetary Union’ funded by the European Union’s Horizon 2020 Program, grant agreement N^o 649396.

[†]Faculty of Economics & Queens’ College, University of Cambridge, and CEPR. Email: cfb46@cam.ac.uk

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

1 Introduction

1.1 Overview

Time inconsistency is an endemic problem in the macroeconomic policy literature. Whether monetary, taxation or social insurance policy, very few meaningful questions can be answered without encountering it in some form. It arises whenever policy must be designed for environments where expectations of future outcomes affect agents' current actions. This dependence provides an incentive to make promises about future policy that it will not be optimal to keep. As a consequence, the 'best' choice of policy instruments for a given time period depends on when this choice is being assessed – is it best *ex-ante*, or contemporaneously? The implied inconsistency in optimal choice was formalised by Kydland and Prescott (1977), and its consequences have been widely studied by macroeconomists ever since.

By definition, time inconsistency means that it is not possible to choose a dynamic allocation that will be optimal from the perspective of every time period in succession. A plan that is optimal initially will not be optimal to continue with. The conventional response to this in the normative policy literature is to surrender the principle of *successive* optimality, and focus on selections that are best from the perspective of the initial time period only. This has commonly come to be known as 'Ramsey' policy design, following the foundational contribution to optimal tax design of Frank Ramsey (1927). It is a method that has been widely applied in many different policy environments.

An alternative approach, comparatively underexplored, is to surrender the principle of optimality, and ask whether there exist weaker normative criteria that *can* be time-consistently satisfied by some dynamic plan. That is, if no policy is best from the perspective of every period, might there nonetheless be options that always remain tolerably good? This is the basic problem that our paper investigates.

Note that this is different from the widely-studied positive question: *What is the equilibrium outcome of discretionary policy choice?* A discretionary outcome is commonly considered a normative failure, implying a lower welfare level in every period than could be attained through a feasible commitment. This is the well-known 'rules beat discretion' result.

Our analysis departs from this positive approach in the *equilibrium* concept used: we assume that policy precommitment is possible. Given this, we differ from the Ramsey approach in the *solution* concept used. In the set of feasible commitments, we seek a policy that exhibits appealing normative properties consistently through time; the Ramsey approach seeks a choice that is optimal for just one period – the first.

Clearly the strength of our approach will rest on what exactly is meant by 'appealing normative properties' here. The analytical device we use to formalise this is the idea

of dominated selections. Even when a general choice problem is subject to time inconsistency, some policy comparisons may be viewed as less contentious than others. For instance, it might be possible to isolate a subset of the available options, and find that for choice in this subset alone, no time inconsistency problem even exists. If this is true, a sub-optimal choice in the restricted subset is surely not desirable for the problem as a whole. Alternatively, it may be that in a pair of feasible alternatives, one option is preferred to the other at every current and future point in time. In this case a Pareto criterion, applied through time, can rule out the inferior choice.

Our paper formalises this reasoning. We endow the space of feasible commitments with an incomplete ‘dominance’ ordering at any given point in time. Where it exists, this ordering always agrees with the policymaker’s preferences, but like the Pareto principle it will leave many pairs of options unranked. Its construction is based on the principles outlined in the previous paragraph. First, in time-consistent subdomains, standard choice is assumed to apply. Second, where the Pareto principle can be applied over time, it is. The idea is that these represent relatively uncontroversial choice principles, even in the wider context of time inconsistency.

Since the dominance ordering agrees with standard policy preferences wherever it exists, the resulting set of undominated allocations will be larger than – and contain – the more exclusive set of optimal choices in each period. Time-consistent membership of the undominated set may thus be possible where time-consistent membership of the optimal set is not. This is the basic normative argument that we pursue.

Given this approach, our main analytical contribution is to characterise necessary and sufficient conditions for policies to belong to the undominated set in every time period. These are the ‘time-consistently undominated policies’ to which the paper’s title refers. We apply these characterisation results to a number of textbook examples, highlighting the differences relative to Ramsey choice in particular. In a version of the Judd (1985) capital tax problem, time-consistently undominated capital taxes satisfy an intuitively simple efficiency-equity trade-off in all periods, and are generally positive. In a social insurance problem with one-sided limited commitment, time-consistently undominated policy involves a stable consumption distribution, with a progressive effective marginal savings tax. In a social insurance problem with asymmetric information, à la Atkeson and Lucas (1992), time-consistently undominated policy induces a stable consumption distribution where Ramsey policy implies an immiseration result.

Central to the general characterisation is a novel set of restrictions on the Lagrange multipliers that attach to dynamic promise-keeping constraints in each period. These multipliers are well-studied objects in the literature on Kydland and Prescott problems, following the work of Marcet and Marimon (1998, 2017). Intuitively they summarise the way that policy balances the prior value of keeping promises with the contemporaneous cost. Ramsey policy requires the multipliers to be highly persistent through time, mean-

ing that the demands of past promises ultimately come to dominate policy choice. Under time-consistently undominated policy the multipliers instead exhibit gradual decay, at a rate that coincides approximately with the policymaker's discount factor. This has significant implications for the character of policy, particularly in the long run.

1.2 Why study this problem?

Our motivation for investigating time-consistent normative solution concepts derives principally from unease expressed in the literature about the properties of Ramsey policy. This remains the main benchmark when generating policy advice, but at least three distinct features make its suitability for practical recommendations questionable.

The first issue relates to the arbitrariness of date-contingent choice. Under Ramsey policy, the optimal instrument choice varies systematically in the amount of time that has elapsed since the initial optimisation period – ‘date zero’. Section 2 provides a simple example. This time variation occurs independently of any evolution in underlying economic variables. A number of authors have argued that such a feature is either undesirable, implausible, or both. As Svensson (1999) put it, “What is special about date zero?”. This view has been particularly prominent in the New Keynesian monetary policy literature, where it prompted Woodford (2003) to develop the widely-applied ‘timeless perspective’ approach to policy design. Though the approach we recommend ultimately differs from Woodford’s, his search for “a systematic *decision procedure* in the light of which ... current actions are always to be justified” is precisely our focus.¹

A second issue with the Ramsey approach relates to its long-run dynamics. There are a number of settings in which the long-run outcomes of a Ramsey-optimal plan can be extremely *undesirable* in isolation. In many dynamic asymmetric information settings, for instance, it may be Ramsey-optimal to drive the consumption of almost all agents to zero as time progresses – even though the policymaker is utilitarian. An example based on Atkeson and Lucas (1992) is given in Section 9.3 below. The deeper problem is that an optimal choice for date zero need not exhibit any clear desirability properties when reassessed at a later point in time. A time-consistent normative choice technique can overcome this by design.

A third feature of Ramsey policy that may be problematic is its relative inflexibility. A Ramsey plan is defined as a set of instrument choices that are optimal from the perspective of date zero. This is crucially dependent on the model of the economy that is adopted in date zero. In practice every model of the economy will come to be updated and improved, in ways that cannot easily be foreseen. How the Ramsey plan should be affected when this occurs is a very difficult problem. Full reoptimisation, treating the current period as a new ‘date zero’, could be viewed as a violation of the past commitment; but retaining the

¹Woodford (2003), § 7.1, p. 474.

existing plan is surely suboptimal. There is no easy intermediate position. Our approach can again overcome this issue. It allows the appropriateness of a policy to be assessed on a rolling basis, without any dependence on past perspectives to motivate choice.

Though we find these arguments interesting and forceful, we also stress that their validity is not our principal concern. It is clear that reasonable doubts *can* exist about the appropriateness of Ramsey policy in certain settings. So long as this is true, it makes sense as a practical matter to investigate normative alternatives.

1.3 Related literature

1.3.1 Commitment, discretion and rules

Since the seminal contribution of Kydland and Prescott (1977), a vast number of papers have engaged with the general problem of time inconsistency – both from a normative and a positive perspective. With the exception of the New Keynesian literature, discussed below, the dominant *normative* focus has been on Ramsey policy, with significant innovations over the years in its characterisation and computation. The work on dynamic games by Abreu, Pearce and Stachetti (1990), and on recursive saddle-point problems by Marcet and Marimon (1998, 2017) has provided alternative devices for representing the Ramsey problem in recursive form.² Our characterisation results, below, are stated in terms of the promise multipliers whose use Marcet and Marimon popularised, and are easiest to interpret by comparison with their work.

The positive literature on time inconsistency considers the implications for policy and welfare of a lack of commitment. Here there are important differences in the equilibrium concept used. The majority of papers seek Markov-perfect equilibria.³ These allow no scope for promises to bind choice, though strategic incentives to influence future decisions can affect the choice of endogenous states. Outcomes are generally inefficient, with commitment strategies delivering welfare improvements from the perspective of every time period.⁴

A smaller, though highly influential, literature focuses on history-contingent reputational equilibria.⁵ This ‘sustainable plans’ approach characterises the set of policies that

²Though Abreu, Pearce and Stachetti (1990) wrote on dynamic games, there have been many applications of their work in the macroeconomics literature, including Kocherlakota (1996a), Chang (1998) and Phelan and Stachetti (2001).

³Examples include Klein and Ríos-Rull (2003), Ortigueira (2006), Ellison and Rankin (2007), Klein, Krusell and Ríos-Rull (2008), Diaz-Giménez, Giovannetti, Marimon and Teles (2008), Martin (2009), Blake and Kirsanova (2012), Reis (2013), Niemann, Pichler and Sorger (2013), Bianchi and Mendoza (2013), and Debortoli, Nunes and Yared (2017).

⁴A related branch of work is the ‘loose commitment’ approach developed by Debortoli and Nunes (2010). This sits between the positive and normative branches of the literature, analysing the outcomes of optimal policy problems when reoptimisation is known to take place at random intervals through time.

⁵Chari and Kehoe (1990) and Atkeson (1991) were pioneering early papers. More recent work of this kind in the social insurance literature includes Sleet and Yeltekin (2006), Sleet and Yeltekin (2008), Acemoglu, Golosov and Tsyvinski (2010), Farhi, Sleet, Werning and Yeltekin (2012) and Golosov and

can be supported by appropriate trigger strategies in an infinite horizon. The threat of reversion to an inferior equilibrium can allow some promises to be kept, though the Ramsey strategy is usually not attainable. A common feature of this literature is indeterminacy: the set of sustainable equilibria is large, though – mirroring Ramsey policy – it is common to focus on the best sustainable equilibria from the perspective of the initial time period.

The variant on this literature that comes closest to our work is Kocherlakota (1996b), who introduces a refinement that he dubs *reconsideration-proofness* to the problem of finding a sustainable plan. Developed in a purely stationary environment, this recommends selecting an equilibrium that is best, subject to the assumption that future policymakers will be allowed to select in exactly the same manner. This naturally leads to the best *constant* choice over time. This exactly coincides with our symmetric time-consistently undominated policy in examples without state variables, though it is not directly applicable to models with states.

1.3.2 The timeless perspective

The problem of finding a time-consistent normative solution concept in Kydland and Prescott problems has been most directly framed in the New Keynesian literature. The ‘timeless perspective’ method proposed by Woodford (1999, 2003) recommends implementing in all periods a policy rule that is consistent with the long-run outcome under Ramsey policy. This method remains commonly applied across a range of problems in monetary policy design, particularly in linear-quadratic environments.⁶

Our results sound a note of caution about the timeless perspective. We show that the long-run continuation of Ramsey policy can generically be Pareto-dominated by alternative feasible selections. This makes the justification for choosing it appear weak.⁷ This is particularly evident in the example of Section 2, where the timeless perspective policy would select a constant inflation-output combination that is strictly inferior to alternative feasible constant policies.

1.3.3 Variable social discounting

The immiseration result is commonly regarded as a troubling conclusion *per se*, and work by Phelan (2006) and Farhi and Werning (2007) investigates options for overcoming it. Like our paper, the approach of these authors is explicitly normative, with the assumption of a perfect commitment device. Unlike our paper, the essential strategy that Phelan

Iovino (2014).

⁶Recent papers making use of it include Giannoni and Woodford (2017), Armenter (2017), Cúrdia and Woodford (2016), Engel (2014), Benigno and Paciello (2014), Adam and Woodford (2012), Benigno and Woodford (2012), and Corsetti, Dedola and Leduc (2010).

⁷The desirability of the timeless approach has been already questioned in the context of a linear-quadratic New Keynesian problem by Blake (2001). See also Damjanovic, Damjanovic and Nolan (2008).

(2006) and Farhi and Werning (2007) propose is to raise the societal discount factor. This is justified on first principles as identifying an alternative position on the intergenerational Pareto frontier.

There is a long tradition in economic policy design, dating at least to Ramsey (1928), that recommends a higher societal discount factor relative to private-sector preferences. Whether this is appropriate or not is a deeply contentious question, and we do not propose to resolve it here. We note simply that it implies a more substantial change to the principles of policy design than our paper. Our method is deliberately designed to preserve standard choice in time-consistent environments. As the example of Section 9.3 shows, it is possible to overcome the immiseration result just by amending choice principles for the time-inconsistent aspects of a problem.

1.4 Paper outline

The paper proceeds as follows. Section 2 outlines a simple linear-quadratic problem that further illustrates the motivation for what we do. Section 3 presents a general problem that we use to develop the main ideas, and discusses some key assumptions. Sections 4 and 5 describe, in turn, the dominance ordering that we place on the space of feasible allocations, and how choice can be conducted in light of this ordering. Section 6 shows that this choice problem can be divided into a two-stage procedure, with a time-consistent ‘inner’ problem that takes promises as given, and a time-inconsistent ‘outer’ problem that is concerned with the choice of promises. This is a crucial step in operationalising our approach.

Section 7 provides necessary and sufficient conditions for policies to be time-consistently undominated, and shows that conventional normative and positive approaches do not satisfy these. Section 8 shows that time-consistently undominated policies have a dual interpretation as promise choices that are optimal for *every* period along *one* choice dimension. This is used to add an appealing symmetry refinement to our approach, allowing multiplicity to be overcome. Section 9 applies our approach to three textbook settings: a capital tax problem, a social insurance problem with limited commitment, and a dynamic asymmetric information problem. Section 10 concludes.

2 Motivating example

The introductory discussion can be clarified by exploring a simple example. This section explains the problem of normative choice in the context of a linear-quadratic New Keynesian inflation bias problem with no uncertainty.⁸ With just two variables and one

⁸The problem is studied for its simplicity rather than its realism. More detailed foundations for it are discussed in Woodford (2003), § 7.1.

linear constraint, the environment is as simple as possible.

2.1 Setup

Time is discrete, and runs infinitely from some initial period 0. The supply side of the economy in period t is described by a linearised New Keynesian Phillips Curve:

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

where π_t is inflation in period t , y_t is a measure of the output gap, $\mathbb{E}_t$ is a standard expectations operator and $\beta \in (0, 1)$ and $\gamma > 0$ are parameters. Policy choice is assumed to be across output and inflation sequences from 0 onwards, subject only to equation (1). To keep notation compact we denote infinite sequences by bold type with an overbar, with subscripts giving the starting period, so $\bar{\mathbf{y}}_0 := \{y_t\}_{t=0}^\infty$, $\bar{\boldsymbol{\pi}}_s := \{\pi_t\}_{t=s}^\infty$, and so on.

2.2 The feasible set

Any pair $(\bar{\mathbf{y}}_s, \bar{\boldsymbol{\pi}}_s)$ that satisfies (1) for all $t \geq s$ is a **feasible choice** from period s onwards. For all $s \geq 0$, define Ξ as the set of feasible policy sequences from s on:

$$\Xi = \{(\bar{\mathbf{y}}_s, \bar{\boldsymbol{\pi}}_s) : (1) \text{ true for all } t \geq s\}$$

Note that Ξ is time-invariant. A pair of inflation and output sequences that is feasible from s onwards would also be feasible from t onwards.

2.3 Time inconsistency and Ramsey choice

The central policy problem is to make a selection from Ξ . We assume a commitment device, so that every element of Ξ can potentially be chosen. The focus is on the normative properties of alternative selections.

In any given period $s \geq 0$ the policymaker has a complete, rational preference ordering over Ξ , described in the usual way by the objective function W_s :

$$W_s := - \sum_{t=s}^{\infty} \beta^{t-s} [\pi_t^2 + \chi (y_t - y^*)^2] \quad (2)$$

where $y^* > 0$ is an optimal level for the output gap and $\chi > 0$ is a parameter.

Ramsey policy is defined as the selection $(\bar{\mathbf{y}}_0^R, \bar{\boldsymbol{\pi}}_0^R)$ such that W_0 is maximised on Ξ . In this simple linear-quadratic environment it will be unique:

$$(\bar{\mathbf{y}}_0^R, \bar{\boldsymbol{\pi}}_0^R) = \arg \max_{(\bar{\mathbf{y}}_0, \bar{\boldsymbol{\pi}}_0) \in \Xi} W_0$$

This policy is an important and widely-studied benchmark, but it is well known that it is a time-inconsistent selection. It recommends values for π_t and y_t that are positive initially, but tend jointly to zero as time progresses. Since the model is entirely stationary, re-optimising in any period $s > 0$ would imply exactly the same dynamics, but starting from s instead of 0. This means departing from the continuation of the period-zero Ramsey plan. Hence ‘maximise W_s on Ξ ’ is not a time-consistent solution concept.

As is well known, the reason for the inconsistency is that constraint (1) contains the forward-looking term $\mathbb{E}_t \pi_{t+1}$. There is an incentive to make promises about future allocations in order to manage inflation expectations. When the future arises, the justification for keeping these promises has passed.

2.4 Time-consistent choice criteria

Time inconsistency implies that a policy cannot be optimal for every period. It can at best either (a) be optimal from the perspective of just one period, or (b) be desirable in some weaker sense, in every period. Choosing the Ramsey plan $(\bar{y}_0^R, \bar{\pi}_0^R)$ means following the first approach, where period 0 is the date that is privileged. Our aim is to operationalise the second approach.

The simplicity of the present example is helpful. Whatever time-consistent solution concept we ultimately devise, in this stationary, deterministic environment it must deliver a *time-invariant* inflation-output choice.⁹ The class of constant inflation-output combinations is easy to investigate here, and provides useful insights that will later generalise.

Formally, we can define the set of feasible *constant* policy options as $\Xi^c \subset \Xi$:

$$\Xi^c := \{(\bar{y}_0, \bar{\pi}_0) \in \Xi : (y_t, \pi_t) = (y_s, \pi_s) \text{ for all } t, s \geq 0\}$$

There is a unique choice that maximises W_s on Ξ^c for all s , which we label $(\bar{y}_0^c, \bar{\pi}_0^c)$ – the optimal constant policy.¹⁰ It would be extremely hard to construct a normative case for any constant choice other than this.

The puzzling aspect of this finding is that $(\bar{y}_0^c, \bar{\pi}_0^c)$ is not related in any obvious way to the main policy benchmarks that exist in the literature. It is neither the long-run outcome from Ramsey policy, nor the time-invariant Markov equilibrium. Figure 1 contrasts optimal constant policy with these outcomes for conventional parameter values.¹¹ The Markov outcome is biased towards excessive inflation, and is clearly Pareto inefficient when considering the preferences of policymakers at differing points in time. More intriguingly, Figure 1 highlights that the *continuation* of Ramsey policy is also inefficient

⁹If this were not true, re-applying the choice criterion in a later period would imply deviating from any earlier selection.

¹⁰It is a simple exercise to show that this is given by $y_t^c = \frac{\chi(1-\beta)^2}{\gamma^2 + \chi(1-\beta)^2} y^*$ and $\pi_t^c = \frac{\chi\gamma(1-\beta)}{\gamma^2 + \chi(1-\beta)^2} y^*$ for all t .

¹¹We assume $\beta = 0.96$, $\gamma = 0.024$, $\chi = 0.048$ and $y^* = 0.05$.




Figure 1: Ramsey, Markov and optimal constant policy

in this sense, once enough time has elapsed. For sufficiently large s , all policymakers from s onwards strictly prefer $(\bar{y}_s^c, \bar{\pi}_s^c)$ to $(\bar{y}_s^R, \bar{\pi}_s^R)$. Note that the limiting outcome of Ramsey policy, with $\pi_t = y_t = 0$, is the ‘timeless perspective’ policy recommended for every period by Woodford (2003). This is clearly inferior to the optimal constant choice.

We can build on this discussion by defining a Pareto dominance across pairs of allocations in this example as follows:

Definition. Policy $(\bar{y}'_s, \bar{\pi}'_s) \in \Xi$ **dominates** the alternative $(\bar{y}''_s, \bar{\pi}''_s) \in \Xi$ in period $s \geq 0$ if there exists $\varepsilon > 0$ such that W_t is higher under $(\bar{y}'_t, \bar{\pi}'_t)$ than $(\bar{y}''_t, \bar{\pi}''_t)$ by at least an amount ε for all $t \geq s$.

A policy $(\bar{y}'_s, \bar{\pi}'_s) \in \Xi$ is **undominated** in period s if there is no alternative in Ξ that dominates it in s . Note that the set of undominated policies in s will always contain the *optimal* policy to implement from s onwards, but it will generally contain many other elements too. Moreover, the optimal choice in s may come to be dominated in continuation in periods subsequent to s . Figure 1 confirms that this is true of the Ramsey policy. We will seek time-consistently undominated policies:

Definition. A policy $(\bar{y}'_0, \bar{\pi}'_0) \in \Xi$ is **time-consistently undominated** if its continuation $(\bar{y}'_s, \bar{\pi}'_s)$ is undominated for all $s \geq 0$.

We have the following result:¹²

Proposition 1. *The optimal constant policy $(\bar{y}_0^c, \bar{\pi}_0^c)$ is time-consistently undominated.*

This is non-trivial, because $(\bar{y}_s^c, \bar{\pi}_s^c)$ is shown to be undominated in the entire set Ξ , not just the restricted set Ξ^c in which it is optimal. Given the chosen definition of dominance, it shows by example that time-consistently undominated policies can exist in environments where time-consistently *optimal* policies do not. Weakening the normative

¹²Proofs of Propositions are collected in Appendix A.




Figure 2: Three time-consistently undominated policies

requirement from ‘optimal’ to ‘undominated’ thus yields a choice criterion that can be asserted in all periods. Moreover, the resulting policy is qualitatively distinct from the main positive and normative benchmarks in the literature.

2.5 Multiplicity

An important qualification to this result is that the optimal constant policy is not the *only* time-consistently undominated selection available. Figure 2 charts two others alongside it. The policy labeled ‘limiting path’ involves a strictly higher inflation rate initially, approaching the optimal constant choice at the limit as time passes. Because the two are equivalent at the limit, the strict Pareto dominance requirement is not met for any $\varepsilon > 0$.¹³ The other policy, labeled ‘fluctuating path’ sees inflation and output follow a two-period cycle, permanently fluctuating about their optimal constant values. This highlights that the multiplicity of time-consistently undominated solutions is not just a ‘transition’ issue. There are time-consistently undominated paths that never converge to the optimal constant solution.

Though it rules out important benchmarks, it is clear that the dominance criterion alone does not deliver a unique time-consistent selection. If uniqueness is desired, some further refinement is necessary. Yet it is also clear that, at least in this example, there is one ‘obvious’ candidate for a refinement – the optimal constant policy. As well as being an order of magnitude simpler, this is the only selection that treats all periods symmetrically.

2.6 Summary

The general lessons from this example can be summarised as follows:

¹³Given this, it is clear that a weaker Pareto criterion would not allow time-consistent choice.

1. In a model without states or shocks, selection from the set of constant policies is a time-consistent choice procedure.
2. The solution to this problem is neither the outcome of a Markovian, discretionary equilibrium, nor the long-run outcome observed under Ramsey policy.
3. The optimal constant policy is time-consistently undominated, according to a strict Pareto criterion.
4. Many other policies are also time-consistently undominated by this criterion, but all of these imply asymmetries in policy choice through time.

The analysis that follows will generalise all four of these insights.

3 General setup

We develop the theory in a general setting that nests a number of the most well-known Kydland and Prescott problems. As above, sequences are written using bold type with an overbar, with subscripts to denote starting period. Superscripts are used to denote the end period of a finite sequence where necessary. Thus $\bar{\mathbf{x}}_{\mathbf{s}} := \{x_t\}_{t \geq s}$, $\bar{\mathbf{x}}_{\mathbf{s}}^r := \{x_t\}_{t=s}^r$, and so on. $(\bar{\mathbf{x}}_{\mathbf{s}}^{r-1*}, \bar{\mathbf{x}}_{\mathbf{r}}')$ denotes the combined sequence $\{x_s^*, x_{s+1}^*, \dots, x_{r-1}^*, x_r', x_{r+1}', \dots\}$.

3.1 Preliminaries

Time is discrete, and runs from period 0 to infinity. We abstract from aggregate risk for simplicity. The framework allows settings with idiosyncratic risk across large populations of agents.

In each period $t \geq 0$ there is a vector of n predetermined ‘state’ variables $x_{t-1} \in X \subset \mathbb{R}^n$, with x_t to be chosen in t , and a vector of m non-predetermined variables $a_t(\sigma) \in A_\sigma \subset \mathbb{R}^m$ defined for all $\sigma \in \Sigma$, where σ is an identifier variable – possibly stochastic – discussed in more detail below, and Σ is the set of possible σ realisations. We define $a_t \in A$ as $\{a_t(\sigma)\}_{\sigma \in \Sigma}$, with $A := \{A_\sigma\}_{\sigma \in \Sigma}$.

The role of σ varies flexibly across examples, but in general it is used to index the set of forward-looking constraints that are of relevance in any given time period. In environments with heterogeneous agents subject to idiosyncratic risk, for instance, each particular $\sigma \in \Sigma$ will correspond to a distinct history of exogenous shocks. Individuals with different shock histories may receive different allocations, and so for each σ a distinct forward-looking restriction may be required. In deterministic environments with multiple forward-looking constraints, σ can be used as a simple index on these constraints.

We assume that Σ is a time-invariant set. In stochastic environments this means that the quantity of information on past shocks across individuals is stationary, not

accumulating over time. This may imply that detailed shock histories for different agents are known even at the start of time, which is a departure from convention in many settings. It would not make a difference for policy results if histories up to period 0 were generated fictitiously, so this information requirement is not a practical impediment to applying our approach.

σ is assumed to follow a Markov process over time, with the conditional probability measure $\Pi(S|\sigma)$ giving the probability of $S \subseteq \Sigma$ in period $t + 1$, given that σ is drawn in t . Where the meaning is obvious, expectations with respect to this measure will be represented by $\mathbb{E}_t$. The conditional measure $\Pi(\cdot|\sigma)$ is assumed to be time-invariant. In addition, there is an unconditional probability measure across the elements in Σ , denoted $\Pi(S)$ for all $S \subset \Sigma$, also independent of time. This satisfies a standard consistency property:

$$\Pi(S) = \int_{\sigma \in \Sigma} \Pi(S|\sigma) d\Pi(\sigma)$$

for all $S \subset \Sigma$.

In environments with idiosyncratic risk, it will often be desirable to link current allocations to individuals' past histories. For this, it is helpful to assume that σ is 'fully revealing' of past type, defined as follows:

Definition. $\sigma' \in \Sigma$ is **fully revealing** of past type if there exist $S \subset \Sigma$ with $\sigma' \in S$ such that there is just one $\sigma \in \Sigma$ with $\Pi(S|\sigma) > 0$.

Assumption 1. For all $\sigma \in \Sigma$, σ is fully revealing of past type.

This assumption, combined with the time-invariance of Σ , implies that in many examples of interest σ will correspond to a complete infinite sequence of past shock draws.

The problem in period s is to select a sequence of the form $(\bar{\mathbf{x}}_s, \bar{\mathbf{a}}_s) \in \mathcal{X} \times \mathcal{A}$, where $\mathcal{X} \times \mathcal{A}$ is the space of infinite sequences of elements in $X \times A$. $\mathcal{X}$ and $\mathcal{A}$ are taken to be Banach spaces, equipped with a norm $\|\cdot\|$. A generic element of $\mathcal{X} \times \mathcal{A}$ is referred to as an **allocation**. This choice problem will be subject to a set of constraints to be discussed below.

3.2 Social preferences

The set of allocations $\mathcal{X} \times \mathcal{A}$ is ordered in generic period s according to some social preference ranking. This ranking is described by the function W_s :

$$W_s := \sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma \in \Sigma} r(a_t(\sigma), \sigma) d\Pi(\sigma) \quad (3)$$

where $r : A_\sigma \times \Sigma \rightarrow \mathbb{R}$ is a within-period, σ -contingent preference function for period $s \geq 0$, and higher values of W_s correspond to more preferred outcomes.

These preferences are dynamically recursive, and so are not themselves a source of time inconsistency. The assumption that r does not depend on any state variables is a useful normalisation without significant loss of generality. It is always possible to define auxiliary constraints and variables that incorporate this dependence.¹⁴

It is useful to define many concepts directly by reference to the binary preference relation that W_s describes on $\mathcal{X} \times \mathcal{A}$. This will be denoted $\succeq$ for weak preference, with $\succ$ and $\sim$ denoting strict preference and indifference respectively. Thus $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succeq (\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ if W_s is weakly higher under $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$, and so on.

3.3 Constraints

There is an i -dimensional vector of ‘structural’ feasibility restrictions linking the inherited and future state vectors in X , and the current variables in A :

$$g(x_{t-1}, x_t, a_t) \geq 0 \quad (4)$$

where $g : X \times X \times A \rightarrow \mathbb{R}^i$. This must be satisfied for all t . An example would be a simple within-period aggregate resource constraint of the form $Y_t - C_t - I_t \geq 0$, or a capital accumulation equation of the form $K_t \leq (1 - \delta) K_{t-1} + I_t$.

Time inconsistency derives from a set of infinite-horizon ‘forward-looking’ constraints, one for each $\sigma \in \Sigma$. These are generally assumed to take the form:

$$\mathbb{E}_t \left[\sum_{\tau=0}^{\infty} \beta^\tau h(a_{t+\tau}(\sigma_{t+\tau}), \sigma_{t+\tau}) \middle| \sigma_t \right] \geq h^0(a_t(\sigma_t), \sigma_t) \quad (5)$$

where $\sigma_{t+\tau} \in \Sigma$ denotes a τ -period successor history to $\sigma_t \in \Sigma$, $h : A_\sigma \times \Sigma \rightarrow \mathbb{R}$ and $h^0 : A_\sigma \times \Sigma \rightarrow \mathbb{R}$ for all $t \geq 0$. When planning choice in period s , condition (5) must be satisfied for all $\sigma_t \in \Sigma$ at all $t \geq s$. The following assumption can be helpful in guaranteeing the relevance of (5):

Assumption 2. *For all $\sigma_t \in \Sigma$ and $\bar{\mathbf{a}}_{t+1} \in \mathcal{A}$, there exists at least one within-period choice $a_t \in A$ such that (5) is violated.*

Assumption 2 helps to keep the constraint space simple in certain choice problems that follow. It could be dispensed with quite easily, but there are expositional gains from using it, as highlighted below.

¹⁴For instance, in a model that features consumption habits it is possible that the desired preference criterion might take the form $r(c_t - \lambda c_{t-1})$ for some variable c_t and parameter λ . In this case we can define $\tilde{c}_t := c_t - \lambda c_{t-1}$, and use this to suppress the dependence of r on the lagged variable c_{t-1} . The definition of $\tilde{c}_t$ then becomes one of the structural restrictions defining the model.

3.3.1 Discussion of constraint (5)

Some limits to generality are necessary to keep the discussion manageable, but constraint (5) is sufficiently flexible to incorporate many of the canonical settings in which time inconsistency features. As with the objective function r , for simplicity we have assumed that state variables do not enter into h or h^0 . This ensures that the space of allocations consistent with (5) alone will be time-invariant. The infinite upper limit in the summation is slightly restrictive, as it rules out examples where only finite-horizon expectations matter. It is straightforward to extend our analysis to allow for such cases, but we avoid doing this to economise on notation.¹⁵

A more significant limitation of (5) is that it does not easily incorporate incentive-compatibility constraints. Unlike (5), incentive restrictions generally require the right-hand side also to be dependent on future policy choices, as individuals compare promised outcomes under alternative behavioural strategies. A variant that would work for this case is:

$$\mathbb{E}_t \left[\sum_{\tau=0}^{\infty} \beta^{\tau} h(a_{t+\tau}(\sigma_{t+\tau}), \sigma_{t+\tau}) \middle| \sigma_t \right] \geq \mathbb{E}_t \left[\sum_{\tau=0}^{\infty} \beta^{\tau} h(a_{t+\tau}(\tilde{\sigma}_{t+\tau}), \sigma_{t+\tau}) \middle| \sigma_t \right] \quad (6)$$

where $\tilde{\sigma}_{t+\tau}$ can be viewed as an admissible type report τ periods after t , potentially different from the agent's true type. This sort of constraint is central to a number of important environments where Kydland and Prescott problems matter, and we do not wish to neglect it in the treatment. The general discussion is restricted to constraints of type (5) to keep notation manageable, but Appendix C extends the main results to problems with constraint (6), and Section 9.3 in the main text provides an application based on the Atkeson and Lucas (1992) problem.

3.3.2 Equivalent h functions

In most settings the function h will have a clear economic interpretation – the within-period level of utility for an agent, for instance, or within-period expenditure. This interpretation conveys economic information beyond what is mathematically necessary to preserve inequality (5), and it may be useful to compare h across dates and states on the basis of this information. This will be particularly useful when formalising a notion of symmetry through time in policy choice.

Formally, define an admissible **equivalence transform** as a function $\phi : \mathbb{R} \times \mathbb{N} \times \Sigma \rightarrow$

¹⁵In many cases the relevant constraint can be rewritten to match the form of (5) even when it does not initially appear to do so. For instance, the New Keynesian Phillips curve in equation (1) can be solved forward to give:

$$\pi_t = \gamma \mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^{\tau} y_{t+\tau}$$

When the equality is read as a two-sided inequality, this maps directly into (5).

$\mathbb{R}$ that can be permitted to transform the h function in inequality (5) without changing its economic content.¹⁶ As is well known, different forms of comparability imply different admissible transforms. There are two main forms of comparability in h that we will consider. The first, and most widely applicable, is that h is difference comparable.

Definition. The function h is **difference comparable** if all admissible transforms take the form:

$$\phi(h, t, \sigma) = \delta h + \alpha_t(\sigma)$$

where the scalar $\delta \in \mathbb{R}_+$ is common across time and states, but the additive coefficient $\alpha_t(\sigma) \in \mathbb{R}$ can vary in both.

As the name suggests, this form of comparability allows changes to h to be meaningfully compared from one date-state to another. It is well known to be a necessary assumption for utilitarian objectives to have meaning.¹⁷ Many important Kydland and Prescott problems assume weighted utilitarian social objectives, whilst also featuring utility-based forward-looking constraints. In these cases, difference comparability across agents' utility functions is implicit in the choice of social welfare function. Difference comparability is also an appropriate assumption to make when treating linearised models.

An alternative is for h to be ratio comparable.

Definition. The function h is **ratio comparable** if all admissible transforms take the form:

$$\phi(h, t, \sigma) = \delta_t(\sigma) h$$

where $\delta_t(\sigma) \in \mathbb{R}_+$ for all t and σ .

This form of comparability implies that proportional changes to h are independently defined. An example of an h function that is ratio comparable is one that specifies an agent's net expenditure within a given period. Proportional increases in expenditure have meaning irrespective of the numeraire used to define value, and remain unaffected as that numeraire is changed. This sort of function features in the implementability condition for many dynamic Ramsey tax problems.

The comparability properties of h are a primitive feature of the economic environment in any given example, defined as part of the specification of h .

3.4 The feasible set

We denote by $\Xi(x_{s-1})$ the **feasible set** of allocations from period s onwards, given x_{s-1} :

$$\Xi(x_{s-1}) = \{(\bar{\mathbf{x}}_s, \bar{\mathbf{a}}_s) \in (\mathcal{X} \times \mathcal{A}) : (4) \ \& \ (5) \text{ true } \forall \sigma \in \Sigma, \forall t \geq s, \text{ given } x_{s-1}\}$$

¹⁶Alongside this will be auxiliary transforms to the h^0 function and the discount factor, so that the mathematical structure of inequality (5) is preserved.

¹⁷See, for instance, Roberts (1980).

Any chosen allocation from period s onwards must be drawn from this set.

It is also convenient to specify in isolation the set of allocations that is consistent with the structural constraints in (4), and the set of allocations that is consistent with the forward-looking constraints (5). The set of allocations that satisfy (4) for any given x_{s-1} is denoted $\Xi^g(x_{s-1})$:

$$\Xi^g(x_{s-1}) = \{(\bar{\mathbf{x}}_s, \bar{\mathbf{a}}_s) \in (\mathcal{X} \times \mathcal{A}) : (4) \text{ true } \forall t \geq s, \text{ given } x_{s-1}\}$$

Similarly, the set of allocations that satisfy constraints (5) will be denoted Ξ^h :

$$\Xi^h = \{(\bar{\mathbf{x}}_s, \bar{\mathbf{a}}_s) \in (\mathcal{X} \times \mathcal{A}) : (5) \text{ true } \forall \sigma_t \in \Sigma, \forall t \geq s\}$$

This set is independent of the initial state vector, since by assumption these do not feature in constraint (5).

3.4.1 Feasibility and possibility

The analysis will make use of an independence of irrelevant alternatives (IIA) condition in what follows. This requires the universe of ‘irrelevant’ alternatives to be specified. In order that restrictions on the basic space $(\mathcal{X} \times \mathcal{A})$ do not impede the applicability of IIA, we adopt the following technical normalisation when defining the constituent space A :

Assumption 3. (*Normalisation of A*) Let $R^h(\sigma)$, $R^{h^0}(\sigma)$ and $R^r(\sigma)$ denote the ranges of the functions $h(\cdot, \sigma)$, $h^0(\cdot, \sigma)$ and $r(\cdot, \sigma)$ respectively, for any given $\sigma \in \Sigma$. For any three functions $\varrho^h : \Sigma \rightarrow R^h(\sigma)$, $\varrho^{h^0} : \Sigma \rightarrow R^{h^0}(\sigma)$ and $\varrho^r : \Sigma \rightarrow R^r(\sigma)$ there exists an $a \in A$ such that $h(a, \sigma) = \varrho^h(\sigma)$, $h^0(a, \sigma) = \varrho^{h^0}(\sigma)$ and $r(a, \sigma) = \varrho^r(\sigma)$.

In words, any combination of values in the ranges of $h(\cdot, \sigma)$, $h^0(\cdot, \sigma)$ and $r(\cdot, \sigma)$ can be attained by some choice of a in A . To the extent that cross-restrictions rule certain combinations out, these restrictions are normalised to belong to the problem’s constraints. They do not describe the universe of possibilities.

This normalisation is sufficient, but by no means necessary for our purposes. It provides the most general guarantee possible that *non-existence* of irrelevant alternatives will never impede the analysis.

3.5 Structural assumptions

To place structure on the problem, we will impose the following assumptions on the main primitives:

Assumption 4. *The functions r , g , h , and h^0 are continuous and bounded. The spaces $A_\sigma \subset \mathbb{R}^m$ and $X \subset \mathbb{R}^n$ are compact and convex.*

Assumption 5. g is quasi-concave, h is concave, h^0 is convex and r is strictly concave.

Assumption 4 provides essential structure and is imposed throughout. Compactness of A_σ and X is its strongest component, as this implies bounds on the set of possible choices that are unrelated to the problem's feasibility constraints. But without loss we can assume that these bounds are set arbitrarily loosely, and never affect the boundaries of the feasible set $\Xi(x_{s-1})$. Assumption 5 is imposed more selectively, as needed. Quasiconcavity in g ensures that the constraint space $\Xi^g(x_{s-1})$ is always convex, and will be useful for deriving sufficiency statements. The remaining concavity and convexity assumptions are needed to obtain some sufficiency results, and to apply the Lagrange multiplier theorem.

4 Ordering

This section provides an axiomatic description of the dominance ordering that we use.

4.1 Basic approach and rationality properties

The analysis proceeds by placing a pairwise ordering on the time-invariant space Ξ^h . This ordering is denoted $\succeq^{TC}$, and is constructed by reference to two axioms that are defined in this section. $\succeq^{TC}$ will be incomplete on Ξ^h , but where it exists it will always agree with the policymaker's preference ranking $\succeq$. This immediately conveys certain basic rationality properties on $\succeq^{TC}$, such as the absence of any cycles in strict preference. Reflexivity of $\succeq^{TC}$ will follow from the axioms, but transitivity is not imposed.¹⁸ The axiomatisation constructs the strict and indifference orderings $\succ^{TC}$ and $\sim^{TC}$ directly, with $\succeq^{TC}$ meaning that either $\succ^{TC}$ or $\sim^{TC}$ is true axiomatically.

Defining $\succeq^{TC}$ on Ξ^h means that the ordering can differ as the problem's forward-looking constraints change, even for the same basic space $\mathcal{X} \times \mathcal{A}$. This reflects the centrality of the time-inconsistency problem to the construction of $\succeq^{TC}$. By construction, $\succeq^{TC}$ will be invariant to the feasibility restrictions that make up $\Xi^g(x_{s-1})$.

4.2 Axiom 1: Constraint-based comparisons

The first axiom is based on isolating restricted subsets of Ξ^h where choice is known to be time-consistent. If the policy problem were restricted to these subsets alone, standard choice techniques could apply without impediment. A standard independence argument then implies that choice for the wider problem should not recover a selection that is inferior within a subset of this kind. The first axiom ensures this property.

¹⁸That is, $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succeq^{TC} (\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ and $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \succeq^{TC} (\bar{\mathbf{x}}'''_s, \bar{\mathbf{a}}'''_s)$ need not imply $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succeq^{TC} (\bar{\mathbf{x}}'''_s, \bar{\mathbf{a}}'''_s)$, since $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ and $(\bar{\mathbf{x}}'''_s, \bar{\mathbf{a}}'''_s)$ may not be ordered.

4.2.1 Time-consistent comparability

The axiom is constructed based on a concept of **time-consistent comparability** across allocations. This provides a formal description of the comparisons for which no time inconsistency problem applies. Intuitively, these are comparisons that can safely be made without any concern that forward-looking constraints could be violated at any horizon. To formalise this, two auxiliary definitions are helpful. The first is the idea of a composite allocation. A composite is constructed by taking the within-period allocations from one or other of a pair of sequences. Formally:

Definition. Fix a pair of allocations $(\bar{x}'_s, \bar{a}'_s), (\bar{x}''_s, \bar{a}''_s) \in \mathcal{X} \times \mathcal{A}$. The allocation $(\bar{x}^*_s, \bar{a}^*_s)$ is a **composite** of $(\bar{x}'_s, \bar{a}'_s)$ and $(\bar{x}''_s, \bar{a}''_s)$ iff for all $t \geq s$, $(x_t^*, a_t^*) \in \{(x'_t, a'_t), (x''_t, a''_t)\}$.

The second definition is of a complete set of allocations.

Definition. The set of allocations $\mathcal{T}_s \subset \mathcal{X} \times \mathcal{A}$ is **complete** iff for every pair $(\bar{x}'_s, \bar{a}'_s), (\bar{x}''_s, \bar{a}''_s) \in \mathcal{T}_s$, every composite of $(\bar{x}'_s, \bar{a}'_s)$ and $(\bar{x}''_s, \bar{a}''_s)$ also belongs to $\mathcal{T}_s$.

The s subscript on $\mathcal{T}_s$ denotes the starting period for the sequences contained within this set. $\mathcal{T}_t$ will then be used to denote the set of continuations of sequences in $\mathcal{T}_s$, for $t > s$, and so on.

Completeness in the set of options from s onwards guarantees time consistency in future choice. Formally:

Proposition 2. Fix $x_{s-1} \in X$. For any complete set of allocations $\mathcal{T}_s \subseteq \Xi^h$, if $(\bar{x}^*_s, \bar{a}^*_s) \in \arg \max_{\mathcal{T}_s \cap \Xi^g(x_{s-1})} W_s$ then $(\bar{x}^*_t, \bar{a}^*_t) \in \arg \max_{\mathcal{T}_t \cap \Xi^g(x_{t-1}^*)} W_t$ for all $t > s$.

Time consistency in *future* choice is one necessary feature of a time-consistent subdomain. A second important requirement is that *prior* forward-looking constraints should not be violated by alternative selections in the initial period s . Both of these requirements are included in the following formal definition of time-consistent comparability.

Definition. The set of allocations $\mathcal{T}_s \subset \Xi^h$ is **time-consistently comparable** to the allocation $(\bar{x}'_s, \bar{a}'_s)$ iff $(\bar{x}'_s, \bar{a}'_s) \in \mathcal{T}_s$ and:

1. $\mathcal{T}_s$ is complete, and
2. For all $t \geq s$, $r > 0$ and $(\bar{x}_{t-r}^{t-1}, \bar{a}_{t-r}^{t-1}) \in (X \times A)^r$, if $((\bar{x}_{t-r}^{t-1}, \bar{x}'_t), (\bar{a}_{t-r}^{t-1}, \bar{a}'_t)) \in \Xi^h$ then $((\bar{x}_{t-r}^{t-1}, \bar{x}''_t), (\bar{a}_{t-r}^{t-1}, \bar{a}''_t)) \in \Xi^h$ for all $(\bar{x}''_t, \bar{a}''_t) \in \mathcal{T}_t$.

The second condition here states that if $(\bar{x}'_t, \bar{a}'_t)$ could be consistent with a certain sequence of outcomes prior to t , then any alternative continuation in $\mathcal{T}_t$ must also be consistent with this sequence.

The most trivial example of a set $\mathcal{T}_s$ that is time-consistently comparable to $(\bar{x}'_s, \bar{a}'_s)$ is the singleton set containing $(\bar{x}'_s, \bar{a}'_s)$ alone. Though not a particularly interesting case, this confirms the basic possibility of satisfying the definition.

4.2.2 Imposing an ordering

By construction it is clear that if $\mathcal{T}_s$ is time-consistently comparable to the allocation $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$, there will be no time inconsistency problem associated with relative comparisons between $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ and other members of $\mathcal{T}_s$. Where there is no time inconsistency, we have no reason to depart from standard choice principles. This motivates the following axiom:

Axiom 1. (*Constraint dominance*) Let $\mathcal{T}_s \subseteq \Xi^h$ be time-consistently comparable to $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in \Xi^h$. Then for all $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \in \mathcal{T}_s$:

1. $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \succ (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ implies $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \succ^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$
2. $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \sim (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ implies $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \sim^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$.

If $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \succ^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ holds by application of Axiom 1, we say that $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ **constraint-dominates** $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$.

Note that the converse $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succ (\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ is not assumed to imply $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succ^{TC} (\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$. The reason for this is that time-consistent comparability only ensures that $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ is consistent with all past constraints that $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ satisfies, and not necessarily vice-versa. A preference for $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ over $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ in period s may coincide with allocation $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ delivering on a tougher set of past promises. If true, it may not be a time-consistent ranking.

4.3 Axiom 2: Preference-based comparisons

The second axiom is based on preference rankings rather than constraint spaces. Within the set of allocations Ξ^h , there will commonly exist pairs for which the within-period ordering $\succeq$ is in agreement through time. A simple example would be any pair of constant allocations in the inflation bias problem of Section 2. Whenever a subset of options has this property, a conventional Pareto principle can justify $\succeq^{TC}$ coinciding with $\succeq$. The second axiom formalises this.

4.3.1 Time-invariant feasibility

A Pareto principle can be applied whenever policy preferences between a pair of allocations, viewed in continuation, are unchanging through time. A complication in applying this idea is that the comparative feasibility of the two alternatives may not be stable, due to the evolution of state variables. For instance, in period s both $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ and $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ may be feasible continuations, but in period $t > s$ $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t)$ may not be feasible, given x'_{t-1} . Depending on the model's structural constraints, this can occur whenever $x'_{t-1} \neq x''_{t-1}$. It makes the application of a Pareto criterion between the two sequences difficult, because $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ and $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t)$ can only be compared in period t under the assumption of

a varying history. But the ordering $\succeq$ only describes preferences across continuations, holding constant past outcomes.¹⁹ It does not necessarily provide a full description of preferences across past decisions, and we do not want the Pareto criterion to be based on an assumption that it does.

For this reason we restrict the definition of dominance to pairwise comparisons that can be made without varying the sequence of state variables. Formally:

Definition. Allocations $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ and $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ in Ξ^h are **preference-comparable** in period s if $\bar{\mathbf{x}}'_s = \bar{\mathbf{x}}''_s$.

Since the feasibility of a sequence in s implies the feasibility of its continuation in $t > s$, the following is immediate:

Remark. If $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ and $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ are preference-comparable and both belong to $\Xi(x_{s-1})$ for some x_{s-1} , then both $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ and $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t)$ belong to $\Xi(x_{t-1})$ where $x_{t-1} = x'_{t-1} = x''_{t-1}$.

That is, the feasibility of a pair of preference-comparable allocations in s implies that both remain feasible in t , under the assumption that one or other of these allocations was pursued up to $t - 1$. Pairwise rankings across continuations will be well defined in every period, given that one or other option has been chosen to date. The Pareto principle will thus be straightforward to apply.

4.3.2 Preference dominance defined

The Pareto principle is asserted in its strong form, so that $\succ^{TC}$ implies the strict ranking $\succ$ holds at all points in time. The example of Section 2 highlighted that weakening this even at the limit as $t \rightarrow \infty$ could make time-consistent choice impossible.

A technical definition of strict preference that will endure in the limit can be achieved by reference to lower contour sets. Let $\mathcal{L}(\bar{\mathbf{a}}_s; \bar{\mathbf{x}}'_s) := \{\bar{\mathbf{a}}'_s \in \mathcal{A} : (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}_s) \succeq (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)\}$ be the lower contour set for the allocation $\bar{\mathbf{a}}_s$ in $\mathcal{A}$ under the ordering $\succeq$, holding constant the sequence of state vectors at $\bar{\mathbf{x}}'_s$. If the norm on $\mathcal{A}$ is denoted by $\|\cdot\|$, then from the definition of a lower contour set we have that $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succ (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}''_s)$ applies if and only if there exists an $\varepsilon > 0$ such that $\|(\bar{\mathbf{a}}'_s - \bar{\mathbf{a}}_s)\| \geq \varepsilon$ for all $\bar{\mathbf{a}}_s \in \mathcal{L}(\bar{\mathbf{a}}'_s; \bar{\mathbf{x}}'_s)$. That is, $\bar{\mathbf{a}}'_s$ is bounded away from the upper contour set of $\bar{\mathbf{a}}''_s$. This can be extended to ensure time-invariant strict preference, including at the limit, by asserting that ε should be uniform over time.

Formally, Axiom 2 on the ordering $\succeq^{TC}$ is the following:

Axiom 2. (Preference dominance) For any pair of preference-comparable allocations $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s), (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}''_s) \in \Xi^h$:

¹⁹Equivalently, the social welfare function W_s that describes $\succeq$ should admit the addition of a separable component that depends only on outcomes prior to s .

1. If there exists an $\varepsilon > 0$ such that for all $t \geq s$ and all $\bar{\mathbf{a}}_t \in \mathcal{L}(\bar{\mathbf{a}}'_t; \bar{\mathbf{x}}'_t)$, $\|(\bar{\mathbf{a}}'_t - \bar{\mathbf{a}}_t)\| \geq \varepsilon$, then $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succ^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}_s)$.
2. If $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t) \sim (\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}''_t)$ for all $t \geq s$, then $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \sim^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}''_s)$.

If $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succ^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}''_s)$ holds by application of Condition 2, we say that $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ **preference-dominates** $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}''_s)$.

Condition 2 ensures that our solution concept will not select the Pareto-inefficient outcomes that can arise as equilibria under discretionary choice.

5 Choice

Orderings are precursors to choice. This section defines the link between the ordering $\succeq^{TC}$ and a robust set of undominated policies in each period. The non-trivial aspect of this is a requirement that chosen policies should be robust to the inclusion of additional ‘irrelevant’ (dominated) alternatives in the feasible set.

5.1 Irrelevant alternatives

The two axioms used to construct the ordering $\succeq^{TC}$ are quite restrictive in their applicability. Axiom 1 only allows comparisons between time-consistently comparable allocations, and Axiom 2 only allows comparisons between allocations that deliver identical paths for the state vector through time. In some settings feasibility can severely restrict the scope to make comparisons of this kind. In extreme cases it may be that the dominance relation cannot be placed on *any* pairs in $\Xi(x_{s-1})$. In these circumstances an expansion of the feasible set could significantly expand the set of dominance comparisons possible, even when the new additions are themselves dominated under $\succ^{TC}$ by other options in $\Xi(x_{s-1})$. This motivates incorporating an ‘independence of irrelevant alternatives’ (IIA) condition into choice.²⁰

In our context an irrelevant choice is an allocation that is dominated under $\succ^{TC}$ by a feasible alternative. More generally, the following definition is used:

Definition. Fix any $x'_{s-1} \in X$. The set $\tilde{\Xi}^g(x'_{s-1}) \supset \Xi^g(x'_{s-1})$ is an **irrelevant extension** of $\Xi^g(x'_{s-1})$ under $\succeq^{TC}$ if for every $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in \tilde{\Xi}^g(x'_{s-1})$ that is **not** in $\Xi^g(x'_{s-1})$:

1. $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in \Xi^h$; and
2. for all $t \geq s$ it is possible to find an allocation $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t) \in \Xi^g(x'_{t-1}) \cap \Xi^h$ such that $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t) \succ^{TC} (\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$.

²⁰Nash (1950) popularised this criterion. Sen (1970) §I*6 contains a useful discussion of it, referring to it as ‘property β ’.

Thus an irrelevant extension is an expansion of the feasible set such that every new allocation is strictly dominated at every point in time by an allocation already in the feasible set.

5.2 A robustly undominated set

To maximise generality we define an undominated set by reference to the largest possible irrelevant extension. Thus denote by $\hat{\Xi}^g(x_{s-1})$ the union of all irrelevant extensions of $\Xi^g(x_{s-1})$ under $\succeq^{TC}$, and let $\hat{\Xi}(x_{s-1}) := \hat{\Xi}^g(x_{s-1}) \cap \Xi^h$. It is immediate that $\hat{\Xi}^g(x_{s-1})$ is the largest possible irrelevant extension of $\Xi^g(x_{s-1})$ under $\succeq^{TC}$. The **undominated set**, $D(x_{s-1})$, is then defined by:

Definition.

$$D(x_{s-1}) = \left\{ (\bar{\mathbf{x}}_s, \bar{\mathbf{a}}_s) \in \Xi(x_{s-1}) : \neg \left[\exists (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in \hat{\Xi}(x_{s-1}) : (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \succ^{TC} (\bar{\mathbf{x}}_s, \bar{\mathbf{a}}_s) \right] \right\}$$

In words, $D(x_{s-1})$ is the set of allocations that is undominated under $\succeq^{TC}$ in every possible irrelevant extension of the constraint set.

6 A two-part problem

This section shows how undominated choices can be analysed via a two-step decomposition of choice into ‘inner’ and ‘outer’ problems. The inner problem is concerned with choice for a given sequence of promises. The outer problem is concerned with the choice of these promises.

6.1 An inner problem: choice given promises

The decomposition makes extensive use of **promise values**. For any allocation $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in \mathcal{X} \times \mathcal{A}$, we will say that this allocation **induces the sequence of promise values** $\bar{\omega}'_s \in \mathcal{W}$, defined elementwise for all $\sigma_{t-1} \in \Sigma$ and all $t \geq s$ by:

$$\omega'_t(\sigma_{t-1}) := \mathbb{E}_{t-1} \left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} h(a'_\tau(\sigma_\tau), \sigma_\tau) \middle| \sigma_{t-1} \right] \quad (7)$$

The space $\mathcal{W}$ to which $\bar{\omega}'_s$ belongs is taken to be a Banach space with norm $\|\cdot\|$. Each $\omega_t(\sigma)$ is defined up to a set of transformations consistent with the definition of h ,²¹ and so $\mathcal{W}$ is likewise.

²¹That is, if h is difference-comparable, $\omega_t(\sigma)$ is defined up to the class of affine transforms with common slope parameter across σ and t , and if h is ratio-comparable, $\omega_t(\sigma)$ is defined up to a scalar multiple (that may vary in σ and t).

Promises play a useful role because they characterise time-consistently comparable sets. Formally:

Proposition 3. *Consider an allocation $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in \Xi^h$, inducing promises $\bar{\omega}'_s$. A complete set $\mathcal{T}_s$ is time-consistently comparable to $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ **if** for all $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \in \mathcal{T}_s$ the following conditions hold for all $\sigma \in \Sigma$:*

$$h(a''_t(\sigma), \sigma) + \beta \omega'_{t+1}(\sigma) \geq h^0(a''_t(\sigma), \sigma) \quad (8)$$

$$\mathbb{E}_{t-1} [h(a''_t(\sigma'), \sigma') + \beta \omega'_{t+1}(\sigma') | \sigma] \geq \omega'_t(\sigma) \quad (9)$$

*When Assumption 2 holds, $\mathcal{T}_s$ is time-consistently comparable to $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ **only if** conditions (8) and (9) hold for all $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \in \mathcal{T}_s$ and all $\sigma \in \Sigma$.*

Assumption 2 is relatively strong, but its role here is mainly to simplify the statement. It could be dispensed with in the ‘only if’ part, but with significant notational cost. Condition (8) will be referred to in what follows as a ‘promise-making constraint’, and (9) as a ‘promise-keeping constraint’.²²

Given Axiom 1, the following problem then becomes central to the analysis:

Problem 1. (Inner Problem)

$$\sup_{(\bar{\mathbf{x}}_s, \bar{\mathbf{a}}_s) \in \Xi^g(x_{s-1})} W_s$$

subject to (8) and (9) for all $t \geq s$ and all $\sigma \in \Sigma$, given $\bar{\omega}'_s \in \mathcal{W}$ and $x_{s-1} \in X$.

The main interest in this problem comes from the following:

Proposition 4. *1. For any $x_{s-1} \in X$, each allocation in the undominated set $D(x_{s-1})$ solves Problem 1 for the promise values that it induces.*

2. Let Assumption 2 hold, and suppose the allocation $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ solves Problem 1 for the promise values that it induces, given $x_{s-1} \in X$. Then no allocation in $\Xi(x_{s-1})$ constraint-dominates $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$.

Thus the Proposition provides conditions under which ‘undominated under Axiom 1’ and ‘solving Problem 1’ are equivalent requirements. For practical purposes part 2 will be more useful than part 1. With enough regularity it is straightforward to find conditions such that allocations solve Problem 1 for the promise values that they induce.

Problem 1 is referred to in what follows as the **inner problem**. By design, it is entirely time-consistent. The **outer problem** is the problem of choosing a sequence of promise values, $\bar{\omega}_s$.

²²Following Abreu, Pearce and Stachetti (1990), it is well known that augmenting the policy design problem with promise-keeping constraints, and treating promises as additional states, allows the Ramsey solution to be recovered using conventional dynamic programming techniques.

6.2 An outer problem: undominated promises

6.2.1 The value function and its derivative

The value of the inner problem can be denoted $V(\bar{\omega}_s; x_{s-1})$, and this function is a useful reference point for analysing the outer choice of promises $\bar{\omega}_s$. It is defined for all $\bar{\omega}_s \in \mathcal{W}$ such that the constraint set for Problem 1 is non-empty, given x_{s-1} . This set is denoted $\Omega(x_{s-1}) \subseteq \mathcal{W}$, and its interior $\mathring{\Omega}(x_{s-1})$. Appendix B establishes conditions for V and Ω to exhibit important regularity properties, notably concavity of V in the promises and convexity of Ω . Assumption 5 is critical to this; it will be harder to treat cases where concavity in h and r and/or convexity in h^0 fail. The appendix also characterises the derivative of V with respect to the promises, where this exists. In general this is given by:

$$\delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s) = \sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma \in \Sigma} \{ \beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] w_{t+1}(\sigma) - \lambda_t^k(\sigma) w_t(\sigma) \} d\Pi(\sigma) \quad (10)$$

where $\bar{\mathbf{w}}_s$ is a vector movement away from $\bar{\omega}_s$, required to have the property that $(\bar{\omega}_s + \alpha \bar{\mathbf{w}}_s) \in \Omega(x_{s-1})$ for all α in a sufficiently small neighbourhood of zero. $\lambda_t^m(\sigma)$ and $\lambda_t^k(\sigma)$ are Lagrange multipliers on (8) and (9) respectively in Problem 1, and σ_- is the predecessor history to σ .

Condition (10) is a standard envelope result, stated for arbitrary derivative vectors. Intuitively, an increase in $\omega_{t+1}(\sigma)$ relaxes the promise-making and promise-keeping constraints in t associated with this history. This accounts for the term $\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]$.²³ Against this, an increase in $\omega_t(\sigma)$ tightens the promise-keeping restriction in period t . This accounts for the term $\lambda_t^k(\sigma)$.

The multipliers $\lambda_t^m(\sigma)$ and $\lambda_t^k(\sigma)$ are important objects in the literature on Kydland and Prescott problems, as highlighted by Marcet and Marimon (1998, 2017). The contrast between our policy recommendations and conventional Ramsey policy is easiest to understand by reference to these objects, and we make extensive use of them below.

6.2.2 Preferences across promises

$V(\bar{\omega}_s; x_{s-1})$ describes a preference ordering over the space $\Omega(x_{s-1})$, particular to the policymaker in period s . This ordering can be denoted $\succeq_{x_{s-1}}^\omega$:

$$\bar{\omega}'_s \succeq_{x_{s-1}}^\omega \bar{\omega}''_s \iff V(\bar{\omega}'_s; x_{s-1}) \geq V(\bar{\omega}''_s; x_{s-1})$$

The next Proposition shows that policies that are undominated according to $\succeq^{TC}$ in *allocation* space are also undominated in *promise* space according to a Pareto criterion

²³Recall that $\omega_{t+1}(\sigma)$ enters into the promise-keeping constraint (9) for the predecessor state σ_- in t , whereas it enters the promise-making constraint (8) for σ itself.

based on $\succeq_{x_{s-1}}^\omega$. To formalise this, a lower contour set $\mathcal{L}^\omega$ is defined in promise space as follows:

$$\mathcal{L}^\omega(\bar{\omega}'_s; x_{s-1}) := \{\bar{\omega}_s \in \mathcal{W} : \bar{\omega}_s \succeq_{x_{s-1}}^\omega \bar{\omega}'_s\}$$

A Pareto criterion based on $\succeq_{x_{s-1}}^\omega$ that mirrors Axiom 2 can be defined on the space of promises as follows:

Definition. Consider the promise sequence $\bar{\omega}'_s$ such that $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ solves Problem 1 for $\bar{\omega}'_s$, given some initial x'_{s-1} . $\bar{\omega}'_s$ is **dominated** by $\bar{\omega}''_s$ if and only if there exists an $\varepsilon > 0$ such that for all $t \geq s$, $\|\bar{\omega}''_t - \bar{\omega}_t\| \geq \varepsilon$ for all $\bar{\omega}_t \in \mathcal{L}^\omega(\bar{\omega}'_t; x'_{t-1})$.

A promise sequence $\bar{\omega}''_s$ dominates $\bar{\omega}'_s$ if at every $t \geq s$, the policymaker would rather switch from $\bar{\omega}'_t$ to $\bar{\omega}''_t$. Note that for this definition, the state x'_{t-1} is the one induced by $\bar{\omega}'_s$.

The interest in this derives from the following Proposition:

Proposition 5. Suppose $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ solves Problem 1 for the promise sequence that it induces, $\bar{\omega}'_s$, given some initial x_{s-1} . Then $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ belongs to $D(x_{s-1})$ if and only if $\bar{\omega}'_s$ is not dominated by any alternative promise sequence.

Thus finding an undominated allocation is equivalent to finding an undominated promise sequence for the outer problem. Undominated promise sequences are relatively easy to identify, and can be characterised directly in terms of the multipliers $\lambda_t^m(\sigma)$ and $\lambda_t^k(\sigma)$.

7 Characterisation results

This section derives necessary and sufficient properties for policies that inhabit the set $D(x_{t-1})$ in *all* time periods. These ‘time-consistently undominated policies’ are the main focus of our paper. The emphasis is on theoretical results that have the greatest generality possible. Section 8 translates these into a more practical method for deriving time-consistently undominated policy, and Section 9 presents direct applications.

7.1 Ramsey policy

We start by recasting the Ramsey problem in terms of the apparatus presented in Section 6. If this problem is posed in period s , then for an initial state vector x_{s-1} it solves:

$$\max_{\bar{\omega}_s \in \Omega(x_{s-1})} V(\bar{\omega}_s; x_{s-1})$$

Consistent with this, and making use of equation (10) above, Ramsey policy will generally require:

$$\lambda_t^k(\sigma) = \lambda_{t-1}^m(\sigma) + \lambda_{t-1}^k(\sigma_-) \quad (11)$$

for Π -almost all $\sigma \in \Sigma$, $t > s$, and:

$$\lambda_s^k(\sigma) = 0 \quad (12)$$

for Π -almost all $\sigma \in \Sigma$. Conditions (11) and (12) are familiar from Marcet and Marimon (1998, 2017). The time inconsistency of the solution is immediate from the difference between (11) and (12).

An important consideration for the current paper is whether, despite its time inconsistency, the Ramsey policy at least remains *undominated* as time progresses, in the sense set out above. It turns out that it does not, at least in environments where the time inconsistency problem prevails indefinitely. Formally:

Proposition 6. *Let $(\bar{x}'_s, \bar{a}'_s)$ solve the Ramsey problem for period s , such that for all $t > s$, $\lambda_t^k(\sigma)$ is bounded above zero for all σ in a positive-measure subset of Σ . Then for all $t > s$, $(\bar{x}'_t, \bar{a}'_t) \notin D(x'_{t-1})$.*

So long as $\lambda_t^k(\sigma)$ remains bounded above zero, time inconsistency remains. In these circumstances, each period's policymaker from $s + 1$ on would accept future promises being less demanding, if inherited promise-keeping constraints were relaxed in return. This delivers a strict improvement in every period, so long as promise-keeping constraints continue to bind indefinitely along the Ramsey path.

7.2 Time-consistently undominated policy: necessity

This subsection provides necessary restrictions on policies that remain in $D(x_{s-1})$ indefinitely. In all cases the Propositions are stated in a way that is independent of equivalent representations of the h function – and hence of promises. This implies slightly different statements depending on whether difference comparability or ratio comparability obtains, and the Propositions allow for this. The essential arguments are identical regardless of the form of comparability.

7.2.1 Long-run dynamics

The first result relates to the long-run evolution of promise multipliers.

Proposition 7. *Suppose that the policy $(\bar{x}'_s, \bar{a}'_s)$ is time-consistently undominated, given some initial x'_{s-1} , and that V is differentiable at the induced promise sequence $\bar{\omega}'_s$. If h is **difference-comparable** then for Π -almost all $\sigma \in \Sigma$, either:*

1. *There is no period τ such that both $[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]$ and $\lambda_t^k(\sigma)$ are bounded above zero for all $t \geq \tau$.*

or :

2. For all $\rho \in (0, 1)$ and all positive scalars K_1 and K_2 , it is possible to find a $\tau \geq s$ and $T > \tau$ such that:

$$K_1 \rho^{r-\tau} < \prod_{t=\tau}^{r-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} < K_2 \left(\frac{1}{\rho}\right)^{r-\tau}$$

for all $r \geq T$.

Identical conditions apply when h is **ratio-comparable**, except that the multiplier objects in part 1 are replaced by $[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \omega_{t+1}(\sigma)$ and $\lambda_t^k(\sigma) \omega_t(\sigma)$ respectively, and the ratio in part 2 is replaced by:

$$\left| \frac{\omega_\tau(\sigma)}{\omega_r(\sigma)} \right| \prod_{t=\tau}^{r-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$$

We focus first on the difference-comparable case. Here the Proposition should be read as a statement about the long-run tendency of the ratio $\frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$. So long as this ratio exists and is bounded above zero in the long run (part 1 of the Proposition), its compounded product from τ onwards must be stable relative to any non-trivial geometric process. The requirement that both $[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]$ and $\lambda_t^k(\sigma)$ are bounded above zero for sufficiently large t reflects a need that promises should not come to be irrelevant to the allocation as time progresses. If the multiplier terms were to converge to zero, the scope to improve welfare by changing promises would clearly be limited. The change to Proposition 7 when the h functions are ratio comparable merely ensures invariance to a change over time to the units in which $\omega_t(\sigma')$ is expressed.

In the event that *convergence* in the multipliers occurs, a far sharper statement is possible:

Corollary 1. Suppose that the policy sequence $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ is time-consistently undominated, and induces multipliers $\lambda_t^k(\sigma)$ and $\lambda_t^m(\sigma)$ that converge to bounded steady-state values $\lambda_{ss}^k(\sigma)$ and $\lambda_{ss}^m(\sigma)$ for all $\sigma \in \Sigma$. Then under difference-comparability:

$$\beta [\lambda_{ss}^m(\sigma) + \lambda_{ss}^k(\sigma_-)] = \lambda_{ss}^k(\sigma) \quad (13)$$

Under ratio comparability the same applies, provided there is additionally convergence in the promise values – something that can always be guaranteed by normalisation. Thus when convergence is assured, time-consistently undominated policy mandates a very simple equality restriction for the promise multipliers, at least in steady state. Again, the contrast with Ramsey policy is worth emphasising. By equation (11), it is immediate that if Ramsey policy converges to a steady state with bounded multipliers, these will satisfy the relationship:

$$\lambda_{ss}^m(\sigma) + \lambda_{ss}^k(\sigma_-) = \lambda_{ss}^k(\sigma)$$

This is inconsistent with (13) whenever the multipliers are non-zero and $\beta < 1$.

7.2.2 Averaging policy preferences over time

Proposition 7 does not place any direct restriction either on the evolution of the promise multipliers from one period to the next, nor on outcomes across different states σ . The next Proposition shows that time-consistently undominated policies must *also* satisfy restrictions along both these dimensions, for ‘almost all’ time periods. The latter concept uses the following:

Definition. Consider an arbitrary vector of variables $z_t \in \mathbb{R}^n$ and an arbitrary function $\phi : Z \rightarrow \mathbb{R}$. For any time period τ and any $\varepsilon > 0$, index by $i \in \{1, \dots, N\}$ the set of periods t in which $|\phi(z_t)| \geq \varepsilon$, with $t(i, \tau, \varepsilon)$ used to denote the time period in which the i th occurrence of this inequality arises subsequent to τ . We will say that the restriction $\phi(z_t) = 0$ is **almost always true** if for all $\varepsilon > 0$ and all τ , either N is finite or:

$$\sup_i [t(i+1, \tau, \varepsilon) - t(i, \tau, \varepsilon)] = \infty$$

We have the following:

Proposition 8. *Suppose that the policy $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ is time-consistently undominated, given some initial $x'_{s-1} \in X$. Then there exists a sequence of scalar values $\{\alpha_t\}_{t \geq s}$ with $\alpha_t \in (0, 1]$ for all $t \geq s$, such that for Π -almost all $\sigma \in \Sigma$, under **difference comparability** the following equality is almost always true:*

$$\alpha_t [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] - \lambda_{t+1}^k(\sigma) = 0 \quad (14)$$

and under **ratio comparability** this becomes:²⁴

$$\alpha_t [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \omega_{t+1}(\sigma) - \lambda_{t+1}^k(\sigma) \omega_{t+1}(\sigma) = 0 \quad (15)$$

Note that the Ramsey optimality conditions (11) and (12) are special cases of (14) with $\alpha_t = 1$ and $\alpha_t = 0$ respectively. The Proposition thus states that time-consistently undominated choice for $t+1$ is a weighted average of two extremes: what a prior policymaker, in t or earlier, would like, and what a policymaker in $t+1$ would like. The relative weight, α_t , is identical across states σ , though it can vary through time.²⁵ In this regard time-consistently undominated policy strikes a balance between the interests of prior and contemporaneous policymakers.

²⁴The $\omega_{t+1}(\sigma)$ in this equation does not cancel, because the definition of ‘almost always true’ is scale-dependent.

²⁵It follows from (13) that $\alpha_t = \beta$ in steady state, if one is reached. But the Proposition does require convergence.

The fact that $\alpha_t > 0$ in Proposition 8 is particularly significant, for two reasons. First, it demonstrates that a Markov discretionary approach to policy design is dominated under the ordering $\succeq^{TC}$, except in trivial cases. Markov policy sets $\lambda_t^k(\sigma) = 0$ for all t and all σ . So long as shadow benefits to making promises endure over time, i.e. $\lambda_t^m(\sigma) > 0$ remains true, this is not compatible with time-consistently undominated choice.

The second implication of $\alpha_t > 0$ relates back to Proposition 5, which established the link between undominated policies and undominated promise sequences. This link was qualified by the requirement that an undominated policy should solve the inner problem *for the promise sequence that it induces*. This is guaranteed only if the promise-keeping constraint is always binding. The Proposition demonstrates that a binding promise-keeping constraint is a generic feature of time-consistently undominated policy, again so long as the benefits from making promises remain positive.

7.3 Time-consistently undominated policy: sufficiency

Propositions 7 and 8 help eliminate important options, notably Markov and Ramsey policy, but for more constructive purposes we need sufficiency results. This section provides general conditions that *guarantee* that a policy never comes to be dominated.

Proposition 9. *Consider a policy $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ that solves Problem 1 for the promise sequence that it induces, $\bar{\omega}'_s$. The continuation of this policy $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ will belong to $D(x'_{t-1})$ for all $t \geq s$ provided the following are true:*

1. *The value function $V(\bar{\omega}_s; x_{s-1})$ is concave in $\bar{\omega}_s$.*
2. (a) *There exist positive scalars $\underline{K}$ and $\bar{K}$ such that for all $\tau \geq s$, $r > \tau$ and $\sigma \in \Sigma$, under **difference comparability**:*

$$\underline{K} \leq \prod_{t=\tau}^{r-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \leq \bar{K} \quad (16)$$

*or, under **ratio comparability**:*

$$\underline{K} \leq \left| \frac{\omega_\tau(\sigma)}{\omega_r(\sigma)} \right| \prod_{t=\tau}^{r-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \leq \bar{K} \quad (17)$$

- (b) *There is a sequence of scalars $\{\alpha_t\}_{t=s}^\infty$, with $\alpha_t \in [\underline{\alpha}, \bar{\alpha}]$ for all t and $0 < \underline{\alpha} \leq \bar{\alpha} < 1$, such that $\lambda_t^m(\sigma)$, $\lambda_t^k(\sigma_-)$ and $\lambda_{t+1}^k(\sigma)$ converge across $\sigma \in \Sigma$ as follows:*

$$\lim_{t \rightarrow \infty} \left[\frac{\lambda_{t+1}^k(\sigma)}{\alpha_t [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \right] = 1 \quad (18)$$

where the rate of convergence is at least linear.

In conventional optimisation problems it is standard for sufficiency conditions to be limited to environments with concave objectives, and part 1 of this Proposition is required for identical reasons. Without it, it would not be possible to reason from local derivative restrictions to a global statement. Conditions for a concave value function are provided in Appendix B.

The second part of the Proposition provides restrictions on the multipliers that are slightly tighter than the necessary restrictions provided in Propositions 7 and 8. A policy is time-consistently undominated provided the compounded ratio of promise multipliers in (16) remains bounded, and provided its multipliers from one period to the next converge to satisfying a common ratio α_t across states – where α_t may be time varying, and is only restricted to lie in some closed range between 0 and 1.

The freedom in α_t permitted under Proposition 9 indicates that it will usually be possible to find many dynamic policies that are time-consistently undominated. This was already demonstrated informally in the inflation bias example of Section 2, with Figure 2 illustrating three alternatives. As the next Section illustrates, an alternative representation of the problem provides an appealing route to resolving this issue.

8 Time consistently undominated policy as a time-consistent optimum

In this section we show that time-consistently *undominated* policies have a parallel interpretation as the time-consistent solution to a restricted *optimisation* problem. This is central in allowing our approach to be operationalised, and to select a unique policy among the multiple time-consistently undominated options.

8.1 A restricted-dimension problem

8.1.1 One-dimensional promise choice

We will consider the problem of choosing a promise sequence $\bar{\omega}_s$ from some restricted-dimensional subspace of $\mathcal{W}$, where this subspace is defined parametrically by reference to a benchmark sequence $\bar{\omega}'_s \in \mathring{\Omega}(x_{s-1})$ and a set of possible vector movements away from $\bar{\omega}'_s$. In order for the analysis to be independent of arbitrary renormalisations, the available vector movements will be defined in a way that is invariant to permissible rescalings of the promise values. Once more, this requires the cases of difference comparability and ratio comparability to be treated distinctly.

Irrespective of the form of comparability, we will define $\bar{\delta}_s$ as an array of ‘slope pa-

rameters' $\delta_t(\sigma)$, with $\delta_t(\sigma) \in [\underline{\delta}, \bar{\delta}]$ for all t and σ , and $0 < \underline{\delta} \leq \bar{\delta} < \infty$:

$$\bar{\delta}_s := \left\{ \left\{ \delta_t(\sigma) \right\}_{\sigma \in \Sigma} \right\}_{t \geq s}$$

θ will denote an array of choice variables $\theta(\sigma) \in \mathbb{R}$, independent of time:

$$\theta := \left\{ \theta(\sigma) \right\}_{\sigma \in \Sigma}$$

Notation is simplified by writing $\theta \bar{\delta}_s$ to denote the array obtained by elementwise multiplication:

$$\theta \bar{\delta}_s := \left\{ \left\{ \theta(\sigma) \delta_t(\sigma) \right\}_{\sigma \in \Sigma} \right\}_{t \geq s}$$

and $\exp \left\{ \theta \bar{\delta}_s \right\}$ to denote the array:

$$\exp \left\{ \theta \bar{\delta}_s \right\} := \left\{ \left\{ \exp \left\{ \theta(\sigma) \delta_t(\sigma) \right\} \right\}_{\sigma \in \Sigma} \right\}_{t \geq s}$$

The realised promise choice $\bar{\omega}_s$ will depend on the chosen value of θ , given $\bar{\delta}_s$ and $\bar{\omega}'_s$. It is written as $\bar{\omega}_s(\theta; \bar{\omega}'_s, \bar{\delta}_s)$:

$$\bar{\omega}_s(\theta; \bar{\omega}'_s, \bar{\delta}_s) := \begin{cases} \bar{\omega}'_s + \theta \bar{\delta}_s & \text{(difference comparability)} \\ \bar{\omega}'_s \exp \left\{ \theta \bar{\delta}_s \right\} & \text{(ratio comparability)} \end{cases}$$

Note that in both of these expressions, varying $\theta(\sigma)$ allows σ -contingent promises to be changed along exactly one dimension for all time periods. A straightforward example of restricted-dimensional choice was seen in the example of Section 2, when studying the set of *constant* inflation-output combinations.

8.1.2 Problem

We will consider the following problem:

Problem 2. (Restricted Promise Choice)

$$\sup_{\theta \in \mathbb{R}^\Sigma} V \left(\bar{\omega}_s(\theta; \bar{\omega}'_s, \bar{\delta}_s) ; x_{s-1} \right)$$

given $\bar{\omega}'_s$ and $\bar{\delta}_s$.

Assuming that it exists, the value of θ that solves this problem is denoted θ^* , with the resulting promise vector $\bar{\omega}_s^* := \bar{\omega}_s(\theta^*; \bar{\omega}'_s, \bar{\delta}_s)$, which is assumed to induce endogenous state vector x_t^* in period $t \geq s$.

Suppose that the solution to Problem 2 induces a promise sequence that belongs to $\mathring{\Omega}(x_{s-1})$, the interior of $\Omega(x_{s-1})$. Then by standard calculus a necessary optimality

condition for Π -almost all σ is:

$$\sum_{r=s}^{\infty} \beta^{r-s} \{ \lambda_r^k(\sigma) \delta_r(\sigma) - \beta [\lambda_r^m(\sigma) + \lambda_r^k(\sigma_-)] \delta_{r+1}(\sigma) \} = 0 \quad (19)$$

for the case of difference comparability and:

$$\sum_{r=s}^{\infty} \beta^{r-s} \{ \lambda_r^k(\sigma) \omega_r(\sigma) \delta_r(\sigma) - \beta [\lambda_r^m(\sigma) + \lambda_r^k(\sigma_-)] \omega_{r+1}(\sigma) \delta_{r+1}(\sigma) \} = 0 \quad (20)$$

for the case of ratio comparability. If the value function is concave in $\bar{\omega}_s$ then these conditions are also sufficient.

8.1.3 Time-consistent choice

Problem 2 is **time-consistent** if the optimal choice θ^* remains the same through time. This is the case of interest. If true, (19) or (20) must hold for *all* possible s , not just as a one-off. This means that the forward sum will cancel, leaving a single within-period restriction that must hold for all $t \geq s$:

$$\lambda_t^k(\sigma) \delta_t(\sigma) - \beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \delta_{t+1}(\sigma) = 0 \quad (21)$$

under difference comparability or:

$$\lambda_t^k(\sigma) \omega_t(\sigma) \delta_t(\sigma) - \beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \omega_{t+1}(\sigma) \delta_{t+1}(\sigma) = 0 \quad (22)$$

with ratio comparability.

Conditions (21) and (22) are within-period cross-restrictions on the multipliers for the promise-keeping and promise-making constraints. By contrast with Ramsey policy, the restriction in the dimensionality of the policy instrument θ is offset by the requirement for choice to be optimal in *every* period, so that a single multiplier restriction for each period still obtains. Promises are chosen optimally for every period along one dimension, rather than being optimal for one period in every dimension.

Suppose that $\bar{\omega}_s$ indeed satisfies these conditions for all t . Then consider the following product ratio under difference comparability:

$$\prod_{t=s}^{r-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} = \frac{\delta_r(\sigma)}{\delta_s(\sigma)}$$

The boundedness restrictions on $\delta_r(\sigma)$ and $\delta_s(\sigma)$ imply that the object on the left-hand side here must be bounded uniformly above 0 and below ∞ in r . Thus $\bar{\omega}_s^*$ will satisfy sufficiency condition 2(a) in Proposition 9, at least for this value of σ . A similar argument

applies under ratio comparability. Summarising as a Proposition:

Proposition 10. *Suppose the value function $V(\bar{\omega}_s; x_{s-1})$ is concave in $\bar{\omega}_s$. Then an allocation $(\bar{x}_s^*, \bar{a}_s^*)$, inducing promises $\bar{\omega}_s^* \in \mathring{\Omega}(x_{s-1})$, satisfies sufficiency condition 2(a) of Proposition 9 if and only if there is a bounded sequence $\bar{\delta}_s$ such that $\bar{\omega}_s^*$ solves Problem 2 recursively, given $\bar{\delta}_s$ and $\bar{\omega}_s' = \bar{\omega}_s^*$.*

This follows from the foregoing discussion, and a formal proof is omitted.

The Proposition does not directly establish that recursive solutions to Problem 2 are time-consistently undominated, as condition 2(a) of Proposition 9 is not enough for this in isolation. The following corollary has more practical applicability for this purpose:

Corollary 2. *Suppose the value function $V(\bar{\omega}_s; x_{s-1})$ is concave in $\bar{\omega}_s$, and that difference comparability applies. If an allocation $(\bar{x}_s^*, \bar{a}_s^*)$, inducing promises $\bar{\omega}_s^*$, solves Problem 2 recursively, given $\bar{\omega}_s' = \bar{\omega}_s^*$ and some $\bar{\delta}_s$, and induces convergence in the intertemporal multipliers to steady-state values $\lambda_{ss}^m(\sigma) > 0$ and $\lambda_{ss}^k(\sigma) > 0$ for Π -almost all $\sigma \in \Sigma$, then $(\bar{x}_s^*, \bar{a}_s^*)$ is time-consistently undominated. The same result applies under ratio comparability if in addition the promises $\omega_t^*(\sigma)$ converge to steady-state values $\omega_{ss}^*(\sigma) \neq 0$ for Π -almost all $\sigma \in \Sigma$.*

Steady-state convergence is not a necessary property for a time-consistently undominated policy, but it is extremely simple to verify when it does arise. Indeed, the most straightforward computational approach to solving for a time-consistently undominated policy will be first to solve for a steady-state allocation, and then to compute convergence to it. This imposes the convergence property directly.

8.2 Symmetry

The unresolved multiplicity in the set of time-consistently undominated policies, first seen in the example of Section 2, is reflected in the number of free parameters that Problem 2 leaves open. Both the ‘intercept’ $\bar{\omega}_s'$ and ‘slope’ $\bar{\delta}_s$ are presently indeterminate. The multiplier restriction (21) or (22) will place one cross restriction on these two choices, but this still leaves one degree of freedom for each date-state.

The example of Section 2 also indicated that a symmetry refinement might resolve this issue. Time-consistent selection from the (restricted-dimensional) set of *constant* policies – an identical problem in each period – delivered the most appealing choice.

More generally, a well-defined version of symmetry is to require that whatever version of Problem 2 recovers the chosen policy, this problem should give each period’s policymaker exactly the *same* control over promises through time. The key step in the argument is that ‘identical promise control’ must be defined in a way that is invariant to admissible renormalisations. This invariance consideration is what ensures uniqueness.

In both of the two comparability cases, it leaves just one possibility for symmetry to have meaning.

Consider first a case of difference comparability. This means that *level* changes in the promises are defined relative to one another, whilst the absolute values of the promises are not. In this case a symmetric version of Problem 2 would require that $\delta_s(\sigma) = \delta_t(\sigma)$ for all $t > s$ and all $\sigma \in \Sigma$. A necessary optimality condition from the resulting time-consistent problem in period t is:

$$\lambda_t^k(\sigma) - \beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] = 0 \quad (23)$$

Equivalently, if ratio comparability holds then *proportional* changes in the promises are all that is defined. A symmetric version of Problem 2 would again require that $\delta_s(\sigma) = \delta_t(\sigma)$ for all $t > s$ and all $\sigma \in \Sigma$, noting that $\delta_t(\sigma)$ is the per-unit proportional change to $\omega_t(\sigma)$ in this case. The following restriction results for all t and σ :

$$\lambda_t^k(\sigma) \omega_t(\sigma) - \beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \omega_{t+1}(\sigma) = 0 \quad (24)$$

Since a symmetric solution imposes the same multiplier restriction each period, it is consistent with a steady state being achieved. Thus sufficiency can be confirmed easily via Corollary 2.

9 Applications

We apply our method to three textbook time-inconsistency problems. These are, first, a capital tax problem in the style of Judd (1985); second, a social insurance problem subject to one-sided limited commitment constraints; and third, a dynamic moral hazard problem in the style of Atkeson and Lucas (1992). Analytical workings are relegated to Appendix D.

9.1 Capital taxation

We consider a variant of the optimal capital tax problem due to Judd (1985), with a balanced budget restriction on government policy.²⁶ This problem has recently received renewed attention through the work of Straub and Werning (2015), who showed that the Ramsey plan may not deliver zero long-run capital taxes, contrary to widespread prior understanding.

²⁶This restriction ensures a forward-looking implementability constraint that must apply in every period, matching our general constraint (5).

9.1.1 Setup

There are two types of agent in equal measure: a worker who supplies labour inelastically and has no access to savings instruments, and a capitalist who does not work. The government's preferences are described by a weighted sum of these agents' lifetime utilities:

$$W_s := \sum_{t=s}^{\infty} \beta^{t-s} [u(c_t^w) + \mu u(c_t^k)] \quad (25)$$

where c_t^w is consumption of the worker in period t , c_t^k is consumption of the capitalist, and $\mu \geq 0$ is the relative Pareto weight on capitalists' welfare.

The period-by-period resource constraint is given by:

$$c_t^w + c_t^k + g_t + k_t \leq f(k_{t-1}) + (1 - \delta) k_{t-1} \quad (26)$$

where k_{t-1} is capital inherited in period t , g_t is an exogenous level of government spending, and the production function f takes as implicit the fixed level of labour supply.²⁷ The government taxes net capital income linearly, with the tax rate denoted τ_t^k . The resulting funds are used to finance government spending and lump-sum transfers to workers, T_t . Since it must run a balanced budget period-by-period, the government's choices must satisfy:

$$(r_t - \delta) \tau_t^k k_{t-1} \geq g_t + T_t \quad (27)$$

where r_t is the rental cost of capital. As is well known,²⁸ this condition can be replaced by an implementability constraint, expressed purely in terms of allocations. We write this in a form consistent with (5) above:

$$u_{c^k,s}(c_s^k + k_s) \leq \sum_{t=s}^{\infty} \beta^{t-s} u_{c^k,t} c_t^k \quad (28)$$

where subscripts denote derivatives in the usual way. Thus the function h^0 here corresponds to $u_{c^k,s}(c_s^k + k_s)$, and h corresponds to $u_{c^k,t} c_t^k$.²⁹ Notice that in general this may not be consistent with concavity in h or convexity in h^0 . This means that the value function will not be guaranteed to be concave.³⁰ We will be able to derive necessary conditions for an optimum, but sufficiency is not guaranteed. As noted by Lucas and Stokey (1983), this problem is shared by conventional Ramsey analysis in the dynamic tax literature. It is not a specific problem with our approach.

²⁷There are the usual constant returns in capital and labour jointly.

²⁸See, for instance, Chari and Kehoe (1999).

²⁹Since k_s is a state variable, strictly it should not be included in the definition of h^0 . Defining an auxiliary variable $\tilde{k}_s$ together with the additional restriction $\tilde{k}_s = k_s$ would allow direct consistency with the general presentation.

³⁰See Appendix B.

Condition (28) can be interpreted as an offer curve for the capitalist in each period, given that the value of private savings must equal the value of future capital. The object $u_{c^k,t}c_t^k$ corresponds to the value of the capitalist's consumption in period 0. Proportional changes to this have meaning independent of normalisations to the price level, whereas level changes do not. Hence the h function is assumed to exhibit **ratio comparability**.

9.1.2 Time-consistently undominated policy

In Appendix D.1 we show that a symmetric time-consistently undominated policy can be characterised in every period by the single condition:

$$k_t \{ \beta \eta_{t+1} [1 + f_{k,t+1} - \delta] - \eta_t \} = c_t^k \{ \eta_t - \mu u_{c^k,t} \} \quad (29)$$

where η_t is the Lagrange multiplier on the resource constraint (26) in period t . Together with (26), (28), and a simple first-order condition with respect to c_t^w (equation (73) in the Appendix), this is sufficient to close the model.

Condition (29) provides an intuitive statement of the trade-off that our policy strikes. The objects in curly brackets can be read as ‘wedges’ relative to a first-best choice. On the left-hand side is the capital wedge, multiplied by the quantity of capital invested in period t . On the right-hand side is the wedge between the shadow cost of resources and the marginal social value of giving income to the capitalist in period t , multiplied by the value of the capitalist's consumption. Intuitively, providing more spending power to the capitalist is desirable to the extent that it boosts savings, and hence reduces the capital wedge. It is undesirable to the extent that it provides resources to an agent whose consumption exceeds the socially desirable level. Condition (29) balances these two concerns.

A Ramsey policy would also incorporate these considerations, but is complicated by an additional desire to tailor intertemporal consumption prices – proportional to $u_{c^k,t}$ – in a way that will be most beneficial from the perspective of period 0. In most cases this gives it an extra degree of dynamic complexity by comparison. Lansing (1999), however, highlighted that the Ramsey problem is substantially simplified when the consumption utility of capitalists is logarithmic. In this case the implementability condition reduces to:

$$\frac{k_s}{c_s} \leq \frac{\beta}{1 - \beta}$$

This is a static restriction, and so the problem is not subject to any time inconsistency problem. This suggests that time-consistently undominated policy should coincide with




Figure 3: Capital tax dynamics: time-consistently undominated policy

Ramsey choice for this case. Indeed, Ramsey policy is easily shown to require:³¹

$$k_t \{ \beta \eta_{t+1} [1 + f_{k,t+1} - \delta] - \eta_t \} = c_t^k \left\{ \eta_t - \mu (c_t^k)^{-1} \right\} \quad (30)$$

The comparison with (29) confirms that time-consistently undominated policy coincides with Ramsey policy when time inconsistency is absent.

Lansing (1999) shows that (30) is consistent with positive long-run capital taxes when μ is set sufficiently low. Straub and Werning (2015) showed that positive long-run capital taxes are a general feature of the problem whenever $\sigma \geq 1$ holds and μ is sufficiently small, with $\sigma = 1$ a threshold case. Our approach provides an alternative generalisation of the $\sigma = 1$ result. It also implies positive long-run capital taxes for small enough μ , but – unlike the Ramsey solutions that Straub and Werning highlight – the simpler dynamics in the promise multipliers relative to the Ramsey case prevent convergence to corner solutions.

Figure 3 illustrates this. It charts the evolution of capital taxes and the capital stock over time, for different values of the initial capital stock, given $\sigma = 2$.³² All variables remain in steady state, conditional on starting there. When the capital stock starts above steady state, capital taxes start above their steady-state values, and likewise taxes are low when the capital stock starts low. Capital income taxes take high values by comparison with conventional results – in the region of 50 per cent. This reflects the fact that the calibration puts zero welfare weight on the capitalist.

³¹This follows from conditions (74) and (75) in Appendix D.1.

³²We assume $\mu = 0$. The production function is Cobb-Douglas with capital share of 0.33, and we set $\beta = 0.96$, $\delta = 0.05$ and $g_t = 0.4$ for all t . The latter corresponds to steady-state government spending of around 25 percent of GDP.

9.2 Limited commitment

We consider a one-sided limited commitment model without savings, in which a continuum of agents receives a stochastic income draw each period. The utilitarian government provides social insurance, subject to a participation constraint.

9.2.1 Setup

Measure $\mu \in [0, 1)$ of the agents are guaranteed to receive a low income y^l in every period, whilst the remaining $(1 - \mu)$ each period receive a high income $y^h > y^l$ with probability p , and y^l with probability $(1 - p)$. The income draws are iid across agents and time, and publicly observable. A utilitarian government seeks to smooth consumption across individuals, subject to ensuring that all individuals are at least as well off as under autarky.

In principle the policymaker has complete information about the entire history of income draws for each agent. However, a sufficient statistic for computing both Ramsey and time-consistently undominated policy is the number of periods elapsed since an agent last received the high income draw, y^h . Thus we let the exogenous stochastic variable $\sigma \in \Sigma$ be defined as the number of periods since a given agent last drew y^h , with $\Sigma := (\mathbb{N} \cup \infty)$.³³ The Markov process governing σ for agents with stochastic incomes is thus:

$$\sigma' = \begin{cases} \sigma + 1 & \text{with prob } (1 - p) \\ 0 & \text{with prob } p \end{cases}$$

Agents with a fixed, low income have $\sigma = \infty$ in all periods.

Given the process determining σ , the utilitarian policymaker ranks continuation allocations from period s onwards according to:

$$W_s := \sum_{t=s}^{\infty} \beta^{t-s} \left[(1 - \mu) \sum_{\sigma=0}^{\infty} (1 - p)^{\sigma} p u(c_t(\sigma)) + \mu u(c_t(\infty)) \right] \quad (31)$$

There is no saving, so the aggregate resource constraint in period t is:

$$(1 - \mu) \sum_{\sigma=0}^{\infty} (1 - p)^{\sigma} p c_t(\sigma) + \mu c_t(\infty) \leq [1 - (1 - \mu)p] y^l + (1 - \mu)p y^h \quad (32)$$

The participation constraints can be written as:

$$\mathbb{E}_s \left[\sum_{t=s}^{\infty} \beta^{t-s} u(c_t(\sigma_t)) \middle| \sigma_s \right] \geq V(\sigma_s) \quad (33)$$

$$\sum_{t=s}^{\infty} \beta^{t-s} u(c_t(\infty)) \geq \frac{u(y^l)}{1 - \beta} \quad (34)$$

³³ $\sigma = \infty$ denotes an agent who has always received y^l .

where σ_t denotes the realisation of σ in period t , given an initial value $\sigma_s \in \mathbb{N}$ in period s , and $V(\sigma)$ is given by:

$$V(\sigma) : = \begin{cases} u(y^h) + \frac{\beta}{1-\beta} [pu(y^h) + (1-p)u(y^l)] & \text{if } \sigma = 0 \\ u(y^l) + \frac{\beta}{1-\beta} [pu(y^h) + (1-p)u(y^l)] & \text{if } \sigma > 0 \end{cases}$$

9.2.2 Ramsey policy

Ramsey policy in this environment has the well-known property that the cross-sectional Pareto weight on agents who are exposed to income risk is non-decreasing over time.³⁴ In initial time periods the promise-making constraints only bind for agents who receive a high income draw. This raises these agents' within-period Pareto weights, and thus their share of consumption. Since aggregate resources are fixed, the effect of this is to reduce the share of consumption going to income-poor agents over time, until eventually even those with permanently low incomes come up against their participation constraint (34). The long-run allocation is one in which income-poor agents are given consumption equal to y^l in perpetuity. These dynamics are charted in Figure 4, for an illustrative calibration.³⁵

9.2.3 Time-consistently undominated policy

Under symmetric time-consistently undominated policy, allocations satisfy the condition:

$$u'(c_t(\sigma))(1 + \beta^\sigma \lambda_t^m(0)) = \eta_t \quad (35)$$

where $\lambda_t^m(0)$ is the multiplier on the promise-making constraint for an agent receiving a high-income shock in t , and η_t is the resource multiplier.

Thus there is geometric decay at rate β in the 'augmented' component of Pareto weights, cross-sectionally, at each point in time. An individual whose current income is high will receive a Pareto weight of $1 + \lambda_t^m(0)$, where $\lambda_t^m(0)$ must be set sufficiently high that this individual wishes to continue participating in the insurance scheme, given the future allocation. An individual whose income was high one period ago (but not today) receives a current Pareto weight of $1 + \beta \lambda_t^m(0)$, and so on. Individuals' allocations depend on their exogenous history, but there is no dependence on past multipliers. The consumption allocation is time-invariant, and the resource multiplier η_t is constant through time. This means low-income individuals are forever able to consume at a level that is elevated above their income – a permanent social security 'safety net'.

³⁴Appendix D.2 provides details.

³⁵Consumption utility is isoelastic, with $\sigma = 1$, and $\beta = 0.96$. We set $y^h = 10$ and $y^l = 1$, with $p = 0.01$ and $\mu = 0.2$.




Figure 4: Consumption dynamics: TCUP vs Ramsey policy

Figure 4 contrasts outcomes under the Ramsey and time-consistently undominated policies. The first panel charts the evolution of consumption over time for an agent who receives a high-income shock, conditional on low income thereafter. It compares time-consistently undominated policy with the dynamics that arise in the Ramsey steady state. As the discussion indicated, consumption decays far more rapidly under TCUP policy than Ramsey, because of the active decay in Pareto weights. A counterpart to this is that consumption levels must be higher under TCUP policy immediately after a shock, so as to preserve participation incentives. The second panel charts the consumption of permanently low-income agents over time. It confirms that Ramsey policy ultimately drives even permanently low earners against their participation constraints. TCUP policy does not.

9.3 Asymmetric information

Finally, we consider a variant of the insurance problem with hidden information due to Atkeson and Lucas (1992).

9.3.1 Setup

The economy receives an aggregate endowment of real income $\bar{Y}$ in each period, and has access to a linear savings technology with gross rate of return $R > 1$. This income must be divided among consumers, who receive unobservable idiosyncratic shocks to their marginal utility of consumption over time. An individual's lifetime utility from period s onwards is:

$$\mathbb{E}_s \left[\sum_{t=s}^{\infty} \beta^{t-s} \theta_t u(c_t) \right] \quad (36)$$

where $\theta_s \in \{\theta^l, \theta^h\}$ is a disturbance to the marginal utility of consumption, with $\theta^h > \theta^l$. This is iid across agents and time, following the process:

$$\theta_t = \begin{cases} \theta^l & \text{with prob } (1 - p) \\ \theta^h & \text{with prob } p \end{cases}$$

We normalise so that $\bar{\theta} := (1 - p)\theta^l + p\theta^h = 1$. The infinite history of an individual's θ draws up to a given period t is denoted by σ_t , with $\sigma_t := \{\theta_t, \theta_{t-1}, \theta_{t-2}, \dots\}$, and Σ is the set of such histories. Time subscripts are omitted from σ where no confusion will arise, and we will use (σ, θ) to denote the history ' σ followed by θ '. Knowledge of this history on the part of the policymaker is assumed in t for all agents, though it would not change the analysis if this were generated fictitiously for periods prior to 0.

The period-by-period resource constraint is:

$$\int_{\sigma \in \Sigma} c_t(\sigma) d\Pi(\sigma) + B_t \leq \bar{Y} + RB_{t-1} \quad (37)$$

where B_t denotes savings in real bonds from t to $t+1$. The utilitarian first-best allocation would imply a higher within-period consumption level for agents who draw θ^h . Since θ is private information, this gives agents with θ^l an incentive to mis-report. Accounting for this, a second-best solution must satisfy an incentive compatibility restriction to guarantee truthful reporting:

$$\mathbb{E}_s \left[\sum_{t=s}^{\infty} \beta^{t-s} \theta_t u(c_t(\sigma_t)) \middle| \sigma_s \right] \geq \mathbb{E}_s \left[\sum_{t=s}^{\infty} \beta^{t-s} \theta_t u(c_t(\tilde{\sigma}_t(\sigma_t))) \middle| \sigma_s \right] \quad (38)$$

where $\tilde{\sigma}_t : \Sigma \rightarrow \Sigma$ denotes an arbitrary reporting strategy for all t , restricted to be self-consistent through time.³⁶ Note that (38) takes the form of constraint (6) rather than (5), for which the general analysis was developed. The extension of the main necessity and sufficiency proofs to this case is straightforward, and developed in Appendix C. It is easy to show that the only binding incentive compatibility constraint will be the restriction that low types should wish to report truthfully, and we proceed on this basis in what follows.

³⁶That is, if σ' is a possible successor history to σ , $\tilde{\sigma}_{t+1}(\sigma')$ must be a possible successor history to $\tilde{\sigma}_t(\sigma)$ for all t .

9.3.2 Ramsey policy

A characteristic feature of Ramsey-optimal choice in this environment is that it satisfies the so-called ‘inverse Euler equation’, given in this case by:³⁷

$$\beta R \mathbb{E}_{t-1} \left[\frac{1}{u_{c,t}(c_t(\sigma_t))} \middle| \sigma_{t-1} \right] = \mathbb{E}_{t-1} \left[\frac{1}{u_{c,t+1}(c_{t+1}(\sigma_{t+1}))} \middle| \sigma_{t-1} \right] \quad (39)$$

With standard preferences we have that $u_{c,t} > 0$, and so both sides of this equation are bounded above zero. In the event that $R \leq \beta^{-1}$, (39) is a supermartingale in the object:

$$\mathbb{E}_{t-1} \left[\frac{1}{u_{c,t}(c_t(\sigma_t))} \middle| \sigma_{t-1} \right]$$

and so this object must converge a.s. to a finite limit. It is possible to show that there are always incentives to induce consumption differentials so long as consumption remains positive, so the implication is that consumption converges to zero for Π -almost all type histories in Σ as time progresses. This is the well-known ‘immiseration’ result, variants of which were first discovered by Green (1987) and Thomas and Worrall (1990).

9.3.3 Time-consistently undominated policy

Under a symmetric time-consistently undominated policy, the equivalent condition to (39) is a period-by-period cross-sectional restriction on individuals’ inverse marginal utilities, relative to a benchmark:

$$\mathbb{E}_{t-1} \left[\frac{1}{u_{c,t}(c_t(\sigma, \theta))} - \frac{1}{\eta_t} \middle| \sigma_- \right] = \beta \mathbb{E}_{t-1} \left[\frac{1}{u_{c,t}(c_t(\sigma_-, \theta))} - \frac{1}{\eta_t} \middle| \sigma_- \right] \quad (40)$$

where σ is a successor history to σ_- , so that (σ, θ) is a history realised two periods after σ_- , and η_t is again the shadow value of resources.

Thus the average value of the inverse marginal utility of consumption in period t , taken across agents who received the given shock history σ_- up to a period $t - s$, converges to $\frac{1}{\eta_t}$ at rate β as s increases. In the event that $R = \beta^{-1}$, η_t will be constant through time, and the outcome will be a time-invariant consumption distribution. The immiseration result no longer applies.

10 Conclusion

Kydland and Prescott problems are environments where it is not possible to choose optimally, all of the time. The challenge for normative policy design is whether to respond to this with a choice that is optimal at just one point in time, or to try to find an

³⁷See Appendix D.3 for derivation.

alternative approach to choice that *can* be implemented in all periods. The purpose of our paper has been to explore the second option. The outcome of this that we propose – time-consistently undominated policy – is particularly interesting because it mandates simple, normatively appealing choices that differ from the Ramsey benchmark both in the short run and the long run. We have shown this both in a general setting, and in a number of textbook examples.

Formally, our analysis is purely normative. It assumes that the policymaker can commit perfectly to a sequence of future choices, and does not analyse the positive question of whether this commitment can be supported in a noncooperative equilibrium.³⁸ But the commitment assumption does raise a positive issue of its own. If it were indeed possible for the policymaker in period 0 to commit to *any* feasible policy, why would they ever fail to select the Ramsey-optimal choice?

A simple answer to this is that in practice governments simply do not appear to design policy rules that exhibit the date-contingent character of Ramsey policy. No central bank, for instance, has an inflation target that depends on the number of years elapsed since the delegation framework was first devised. There appears to be a practical desire to avoid arbitrary time variation in policy, and a theory that enables this formally can only aid macroeconomic policy design.

A more subtle response relates to the connection between the choice procedure and the commitment assumption itself. In reality no society has access to a perfect commitment device *ex-ante*, resistant to all conceivable challenges. Laws can always be repealed, and constitutions amended or rewritten.³⁹ But a commitment may be particularly exposed to challenge if its continuation cannot be justified by reapplying the principles that selected it in the first place. If an optimal policy was appropriate yesterday, why not today? The normative principles that we set out in this paper allow choice that will be robust to this sort of challenge. In itself this may make the commitment assumption far more credible.

References

- [1] Abreu, D., D. Pearce and E. Stacchetti (1990), ‘Toward a Theory of Discounted Repeated Games with Imperfect Monitoring’, *Econometrica*, 58 (5), 1041–1063.
- [2] Acemoglu, D., M. Golosov and A. Tsyvinski (2010), ‘Dynamic Mirrlees Taxation under Political Economy Constraints’, *The Review of Economic Studies*, 77(3), 841–881.

³⁸This is similar to the distinction between cooperative approaches to game theory, in the spirit of Nash (1950) and Shapley (1953), and non-cooperative approaches, following Nash (1951).

³⁹In some cases this flexibility may be desirable, notably if the structure of the economy has changed in a way that was simply unanticipated when the policy was first designed.

- [3] Adam, K., and M. Woodford (2012), ‘Robustly Optimal Monetary Policy in a Micro-founded New Keynesian Model’, *Journal of Monetary Economics*, 59(5), 468–487.
- [4] Armenter, R. (2017), ‘The Perils of Nominal Targets’, *The Review of Economic Studies*, forthcoming.
- [5] Atkeson, A. (1991), ‘International Lending with Moral Hazard and Risk of Repudiation’, *Econometrica*, 59(4), 1069–1089.
- [6] Atkeson, A., and R.E. Lucas Jr. (1992), ‘On Efficient Distribution with Private Information’, *Review of Economic Studies*, 59(3), 427–453.
- [7] Benigno, P., and L. Paciello (2014), ‘Monetary policy, Doubts and Asset Prices’, *Journal of Monetary Economics*, 64, 85–98
- [8] Benigno, P., and M. Woodford (2012), ‘Linear-quadratic Approximation of Optimal Policy Problems’, *Journal of Economic Theory*, 147(1), 1–42.
- [9] Bianchi, J., and E.G. Mendoza (2013), ‘Optimal Time-Consistent Macroprudential Policy’, NBER WP w19704.
- [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] Chamley, C. (1986), ‘Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives’, *Econometrica*, 54(3), 607–622.
- [13] Chang, R. (1998), ‘Credible Monetary Policy with Long-Lived Agents: Recursive Approaches’, *Journal of Economic Theory*, 81, 431–461.
- [14] Chari, V.V., and P.J. Kehoe (1990), ‘Sustainable Plans’, *Journal of Political Economy*, 98, 783–802.
- [15] Chari, V.V., and P.J. Kehoe (1999), ‘Optimal Fiscal and Monetary Policy’, in J.B. Taylor and M. Woodford (eds.) *Handbook of Macroeconomics*, Vol. 1C, Amsterdam: Elsevier.
- [16] Corsetti, G., L. Dedola and S. Leduc (2010), ‘Optimal Monetary Policy in Open Economies’, in B.M. Friedman and M. Woodford (eds.), *Handbook of Monetary Economics*, Vol. 3, Amsterdam: Elsevier.
- [17] Cúrdia, V., and M. Woodford (2016), ‘Credit Frictions and Optimal Monetary Policy’, *Journal of Monetary Economics*, 84, 30–65.

- [18] Damjanovic, T., V. Damjanovic, and C. Nolan. (2008), ‘Unconditionally Optimal Monetary Policy’, *Journal of Monetary Economics*, 55, 491–500.
- [19] Debortoli, D., and R. Nunes (2010) ‘Fiscal Policy under Loose Commitment’, *Journal of Economic Theory*, 145(3), 1005–1032.
- [20] Debortoli, D., R. Nunes and P. Yared (2017), ‘Optimal Time-Consistent Government Debt Maturity’, *Quarterly Journal of Economics*, 132(1), 55–102.
- [21] 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.
- [22] 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.
- [23] Engel, C. (2014), ‘Exchange Rate Stabilization and Welfare’, *Annu. Rev. Econ.*, 6(1), 155–177.
- [24] Farhi, E., and I. Werning (2007), ‘Inequality and Social Discounting’, *Journal of Political Economy*, 115(3), 365–402.
- [25] Farhi, E., C. Sleet, I. Werning and S. Yeltekin (2012), ‘Non-linear Capital Taxation Without Commitment’, *Review of Economic Studies*, 79(4), 1469–1493.
- [26] Giannoni, M.P., and M. Woodford (2017), ‘Optimal Target Criteria for Stabilization Policy’, *Journal of Economic Theory*, 168, 55–106.
- [27] Golosov, M., and L. Iovino (2014), ‘Social Insurance, Information Revelation, And Lack Of Commitment’, NBER WP w20633.
- [28] Green, E.J. (1987), ‘Lending and the Smoothing of Uninsurable Income’, in E.C. Prescott and N. Wallace (eds.) *Contractual Arrangements for Intertemporal Trade*, Minneapolis: University of Minnesota Press.
- [29] Judd, K. (1985), ‘Redistributive Taxation in a Simple Perfect Foresight Model’, *Journal of Public Economics*, 28, 59–83.
- [30] Klein, P., P. Krusell and J.-V. Ríos-Rull (2008), ‘Time-Consistent Public Policy’, *Review of Economic Studies*, 75, 789–808.
- [31] Klein, P., and J.-V. Ríos-Rull (2003), ‘Time-Consistent Optimal Fiscal Policy’, *International Economic Review*, 44(4), 1217–1246.
- [32] Kocherlakota, N.R. (1996a), ‘Implications of Efficient Risk Sharing without Commitment’, *Review of Economic Studies*, 63(4), 595–609.

- [33] Kocherlakota, N.R. (1996b), ‘Reconsideration-Proofness: A Refinement for Infinite Horizon Time Inconsistency’, *Games and Economic Behavior*, 15, 33–54.
- [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] Lucas, R.E., and N.L. Stokey (1983), ‘Optimal Fiscal and Monetary Policy in an Economy Without Capital’, *Journal of Monetary Economics*, 12(1), 55–93.
- [36] Luenberger, D.G. (1969), *Optimization by Vector Space Methods*, New York: Wiley.
- [37] Marcet, A., and R. Marimon (1998), ‘Recursive Contracts’, European University Institute WP ECO 1998 #37.
- [38] Marcet, A., and R. Marimon (2017), ‘Recursive Contracts’, Manuscript.
- [39] Martin, F. (2009), ‘A Positive Theory of Government Debt’, *Review of Economic Dynamics*, 12(4), 608–631.
- [40] Nash, J.F. (1950), ‘The Bargaining Problem’, *Econometrica*, 18(2), 155–162.
- [41] Nash, J.F. (1951), ‘Non-Cooperative Games’, *Annals of Mathematics*, 54(2), 286–295.
- [42] 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.
- [43] Ortigueira, S. (2006), ‘Markov-perfect optimal taxation’, *Review of Economic Dynamics*, 9(1), 153–178.
- [44] Phelan, C. (2006), ‘Opportunity and Social Mobility’, *Review of Economic Studies*, 73(2), 487–505.
- [45] Phelan, C., and E. Stachetti (2001), ‘Sequential Equilibria in a Ramsey Tax Model’, *Econometrica*, 69(6), 1491–1518.
- [46] Ramsey, F.P. (1927), ‘A Contribution to the Theory of Taxation’, *The Economic Journal*, 37(145), 47–61.
- [47] Ramsey, F.P. (1928), ‘A Mathematical Theory of Saving’, *The Economic Journal*, 38(152), 543–559.
- [48] Reis, C. (2013), ‘Taxation Without Commitment’, *Economic Theory*, 52(2), 565–588.
- [49] Roberts, K.W.S. (1980), ‘Interpersonal Comparability and Social Choice Theory’, *Review of Economic Studies*, 47(2), 421–439.

- [50] Sen, A.K. (1970), *Collective Choice and Social Welfare*, Amsterdam: Elsevier.
- [51] Shapley, L.S. (1953), ‘A Value for n-person Games’, *Contributions to the Theory of Games*, 2(28), 307–317.
- [52] Sleet, C., and S. Yeltekin (2006), ‘Credibility and Endogenous Societal Discounting’, *Review of Economic Dynamics*, 9(3), 410–437.
- [53] Sleet, C., and S. Yeltekin (2008), ‘Politically Credible Social Insurance’, *Journal of Monetary Economics*, 55(1), 129–151.
- [54] Straub, L., and I. Werning (2015), ‘Positive Long Run Capital Taxation: Chamley-Judd Revisited’, Working Paper.
- [55] Svensson, L.E.O. (1999), ‘How Should Monetary Policy Be Conducted In An Era Of Price Stability?’, in *New Challenges for Monetary Policy*, Kansas City: Federal Reserve Bank of Kansas City.
- [56] 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.
- [57] 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.
- [58] Woodford, M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.

A Proofs of Propositions

Proof of Proposition 1

As noted in the text, it is simple to show that the optimal constant policy implies the following values for y_t^c and π_t^c for all t :

$$y_t^c = \frac{\chi(1-\beta)^2}{\gamma^2 + \chi(1-\beta)^2} y^* \quad (41)$$

$$\pi_t^c = \frac{\chi\gamma(1-\beta)}{\gamma^2 + \chi(1-\beta)^2} y^* \quad (42)$$

For this policy to be dominated in some period s , there would have to exist an alternative policy $(\bar{y}_s', \bar{\pi}_s')$ such that the loss associated with (the continuation of) this policy is strictly lower in every period from s on. The constraint set is linear and the

loss function is convex, so this in turn implies that a *differential* movement from $(\bar{\mathbf{y}}_s^c, \bar{\pi}_s^c)$ along the vector $[(\bar{\mathbf{y}}_s^c, \bar{\pi}_s^c) - (\bar{\mathbf{y}}_s', \bar{\pi}_s')]$ must be welfare-improving at the margin. Denote the corresponding sequence of derivatives $\left\{ \frac{dy_t}{d\Delta}, \frac{d\pi_t}{d\Delta} \right\}_{t=s}^\infty$, where Δ is a normalisation factor. Since policy choices under both alternatives are bounded for all t , the derivatives must also satisfy a bound: $\left| \frac{d\pi_t}{d\Delta} \right| < \bar{\Pi}$ and $\left| \frac{dy_t}{d\Delta} \right| < \bar{Y}$ for all t and some $\bar{\Pi}$ and $\bar{Y}$ values. Since $(\bar{\mathbf{y}}_s', \bar{\pi}_s')$ is a strict improvement on $(\bar{\mathbf{y}}_s^c, \bar{\pi}_s^c)$ for all $r \geq s$, by definition there must exist some value $\delta > 0$, independent of r , such that the following is true for all $r \geq s$:

$$-\sum_{t=r}^{\infty} \beta^{t-r} \left[\pi_t^c \frac{d\pi_t}{d\Delta} + \chi(y_t^c - y^*) \frac{dy_t}{d\Delta} \right] \geq \delta \quad (43)$$

From the Phillips curve constraint, we know:

$$\frac{dy_t}{d\Delta} = \frac{1}{\gamma} \left[\frac{d\pi_t}{d\Delta} - \beta \frac{d\pi_{t+1}}{d\Delta} \right] \quad (44)$$

Substituting this into inequality (43) gives:

$$\frac{\beta\chi\gamma}{\gamma^2 + \chi(1-\beta)^2} y^* \left\{ \sum_{t=r}^{\infty} \beta^{t-r} \frac{d\pi_t}{d\Delta} - \sum_{t=r+1}^{\infty} \beta^{t-r-1} \frac{d\pi_t}{d\Delta} \right\} \geq \delta \quad (45)$$

Define $D_r := \sum_{t=r}^{\infty} \beta^{t-r} \frac{d\pi_t}{d\Delta}$. Notice that since $\frac{d\pi_t}{d\Delta}$ is uniformly bounded in absolute value for all t , D_r is uniformly bounded in absolute value for all r . But condition (45) can be rewritten as:

$$D_{r+1} \leq D_r - \tilde{\delta} \quad (46)$$

where $\tilde{\delta} := \delta \left[\frac{\beta\chi\gamma}{\gamma^2 + \chi(1-\beta)^2} y^* \right]^{-1} > 0$. Since this must hold for all $r \geq s$, the boundedness of D_r is contradicted.

Proof of Proposition 2

Suppose the alternative allocation $(\bar{\mathbf{x}}_t', \bar{\mathbf{a}}_t') \in \mathcal{T}_t \cap \Xi^g(x_{t-1}^*)$ was strictly preferred to $(\bar{\mathbf{x}}_t^*, \bar{\mathbf{a}}_t^*)$ in period $t > s$ under W_t . Since preferences are recursive, the composite allocation $((\bar{\mathbf{x}}_r^{t-1*}, \bar{\mathbf{x}}_t'), (\bar{\mathbf{a}}_r^{t-1*}, \bar{\mathbf{a}}_t'))$ would then be strictly superior to $(\bar{\mathbf{x}}_s^*, \bar{\mathbf{a}}_s^*)$ from the perspective of period s . But since $\mathcal{T}_s$ is complete, this composite belongs to $\mathcal{T}_s \cap \Xi^g(x_{s-1})$. This contradicts $(\bar{\mathbf{x}}_s^*, \bar{\mathbf{a}}_s^*)$ being in the arg max set in period s .

Proof of Proposition 3

The two conditions required for a time-consistently comparable set are set out in the definition in Section 4.2.1. The ‘if’ part of the Proposition is straightforward. Completeness (part 1 of the definition) is true by assumption, so only condition 2 in the definition

needs to be confirmed. This follows immediately from (9), which guarantees that any past promises that $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ keeps must also be respected by all other policies in the constraint set, for all $t \geq s$.

For ‘only if’, suppose first that an allocation $(\bar{\mathbf{x}}_s^*, \bar{\mathbf{a}}_s^*) \in \mathcal{T}_s$ violated (8) for some σ and t . By the completeness of $\mathcal{T}_s$, we have $((\bar{\mathbf{x}}_s^{t*}, \bar{\mathbf{x}}'_{t+1}), (\bar{\mathbf{a}}_s^{t*}, \bar{\mathbf{a}}'_{t+1})) \in \mathcal{T}_s$. But since $\omega'_{t+1}(\sigma)$ is generated by $(\bar{\mathbf{x}}'_{t+1}, \bar{\mathbf{a}}'_{t+1})$, this allocation violates (5) for the given σ and t . This is inconsistent with $\mathcal{T}_s \subseteq \Xi^h$.

Suppose instead that $(\bar{\mathbf{x}}_s^*, \bar{\mathbf{a}}_s^*) \in \mathcal{T}_s$ violates (8) when $\sigma = \sigma_{t-1}$ in some period t . Assumption 2 together with the continuity of h and h^0 and compactness of A_σ (Assumption 4) implies it will always be possible to find $\tilde{a}_{t-1}$ such that, for this σ_{t-1} realisation:

$$\mathbb{E}_{t-1} \left[h(\tilde{a}_{t-1}(\sigma_{t-1}), \sigma_{t-1}) + \beta \sum_{\tau=0}^{\infty} \beta^\tau h(a'_{t+\tau}(\sigma_{t+\tau}), \sigma_{t+\tau}) \middle| \sigma_{t-1} \right] = h^0(\tilde{a}_{t-1}(\sigma_{t-1}), \sigma_{t-1}) \quad (47)$$

Now consider the composite allocation $((x_t^*, \bar{\mathbf{x}}'_{t+1}), (a_t^*, \bar{\mathbf{a}}'_{t+1})) \in \mathcal{T}_t$. The maintained hypothesis is:

$$\begin{aligned} \mathbb{E}_{t-1} [h(a_t^*(\sigma_t), \sigma_t) + \beta \omega'_{t+1}(\sigma_t) \middle| \sigma_{t-1}] &< \omega'_t(\sigma_t) \\ &= \mathbb{E}_{t-1} \left[\sum_{\tau=0}^{\infty} \beta^\tau h(a'_{t+\tau}(\sigma_{t+\tau}), \sigma_{t+\tau}) \middle| \sigma_{t-1} \right] \end{aligned}$$

Combining this with (47) and using the definition of $\omega'_{t+1}(\sigma_t)$:

$$\begin{aligned} &\mathbb{E}_{t-1} \left\{ h(\tilde{a}_{t-1}(\sigma_{t-1}), \sigma_{t-1}) + \beta \left[h(a_t^*(\sigma_t), \sigma_t) + \sum_{\tau=1}^{\infty} \beta^\tau h(a'_{t+\tau}(\sigma_{t+\tau}), \sigma_{t+\tau}) \right] \middle| \sigma_{t-1} \right\} \\ &< h^0(\tilde{a}_{t-1}(\sigma_{t-1}), \sigma_{t-1}) \end{aligned}$$

Thus the composite allocation violates a prior forward-looking constraint that $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ satisfies. This contradicts condition 2 of the definition of a time-consistently comparable set.

Proof of Proposition 4

For part 1 suppose otherwise, and take an allocation $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in D(x_{s-1})$ that does not solve Problem 1 for the promise values that it induces. Then there is an allocation $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ in the constraint set for Problem 1 with $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \succ (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$. But belonging to the constraint set for Problem 1 implies that $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ satisfies conditions (8) and (9) for the promise values that $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ induces. Thus (by Proposition 3) $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ belongs to a complete set of allocations that is time-consistently comparable to $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$, and so by Condition 1 $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \succ^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$. This contradicts $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in D(x_{s-1})$.

For part 2 it is sufficient to show that any feasible set of allocations that is time-consistently comparable to $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ must belong to the constraint set for Problem 1, given the promise values that $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ induces. This is true under Assumption 2, as Proposition 3 established. Since $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ is an optimal choice for Problem 1, constraint dominance is not possible.

Proof of Proposition 5

Take the ‘if’ part of the claim first, and suppose otherwise – so that $\bar{\omega}'_s$ is undominated, but $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \notin D(x_{s-1})$. That is, there is no alternative promise sequence $\bar{\omega}''_s$ that dominates $\bar{\omega}'_s$ when the initial state vector is x'_{s-1} , but there is an alternative allocation $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \in \tilde{\Xi}^g(x'_{s-1}) \cap \Xi^h$ such that $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \succ^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$, where $\tilde{\Xi}^g(x'_{s-1})$ is an irrelevant extension of $\Xi^g(x'_{s-1})$. Since $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ solves Problem 1 for the promise sequence that it induces, $(\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s) \succ^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ cannot be applying through constraint dominance. Thus we must have $\bar{\mathbf{x}}''_s = \bar{\mathbf{x}}'_s$, and preference dominance applying such that $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t) \succ (\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t)$ for all $t \geq s$, and at the limit as $t \rightarrow \infty$.

Consider the promise sequence that $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ induces, denoted $\bar{\omega}'_s$. If $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in \Xi^g(x'_{s-1})$, then it is immediate that $\bar{\omega}''_s$ dominates $\bar{\omega}'_s$, since a switch to $\bar{\omega}''_t$ can guarantee at least as desirable an outcome as $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ for all $t \geq s$. Thus $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \notin \Xi^g(x'_{s-1})$. But then it follows from the definition of irrelevant extensions that there is a set of alternative allocations $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t) \in \Xi^g(x'_{t-1}) \cap \Xi^h$ for all $t \geq s$ (with the chosen $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t)$ potentially varying in t) such that $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t) \succ^{TC} (\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$. This ordering either applies through constraint dominance or preference dominance. For constraint dominance to apply, by Proposition 4 each $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t)$ must satisfy the constraint set for Problem 1 generated by the promise sequence $\bar{\omega}'_t$, and deliver higher welfare than $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ for all $t \geq s$. Since $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ in turn delivers higher welfare than $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ for all $t \geq s$, including at the limit, it follows that $\bar{\omega}''_t \succ_{x'_{t-1}}^{\omega} \bar{\omega}'_t$ for all $t \geq s$, including at the limit – a contradiction.

The only remaining possibility is that $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t) \succ^{TC} (\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ holds by preference dominance for all $t \geq s$. In this case $\bar{\mathbf{x}}''_s = \bar{\mathbf{x}}'_s$, and it is immediate that the promise sequence that $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ induces, say $\bar{\omega}'''_s$, dominates $\bar{\omega}'_s$. This contradiction establishes the first part of the result.

For the ‘only if’ part, suppose otherwise – so that $\bar{\omega}'_s$ is dominated, but $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s) \in D(x_{s-1})$. That is, there is no alternative allocation $(\bar{\mathbf{x}}^*_s, \bar{\mathbf{a}}^*_s) \in \tilde{\Xi}^g(x'_{s-1}) \cap \Xi^h$ such that $(\bar{\mathbf{x}}^*_s, \bar{\mathbf{a}}^*_s) \succ^{TC} (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$, where $\tilde{\Xi}^g(x'_{s-1})$ is an irrelevant extension of $\Xi^g(x'_{s-1})$, but there is an alternative promise sequence $\bar{\omega}''_s$ that dominates $\bar{\omega}'_s$ when the initial state vector is x'_{s-1} . Since $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ solves Problem 1 for the promise sequence $\bar{\omega}'_s$, this means that for all $t \geq s$ (and at any limit as $t \rightarrow \infty$) there exists a sequence $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t) \in \Xi^g(x'_{t-1}) \cap \Xi^h$ such that $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t) \succ (\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$, with $(\bar{\mathbf{x}}''_t, \bar{\mathbf{a}}''_t)$ satisfying the constraints for Problem 1 when the promise sequence is $\bar{\omega}''_t$. Denote by W''_t the value of the social welfare criterion

when $(\bar{\mathbf{x}}_t'', \bar{\mathbf{a}}_t'')$ is implemented, and W_t' when $(\bar{\mathbf{x}}_t', \bar{\mathbf{a}}_t')$ is implemented. Likewise, r_t' is used as shorthand for $\sum_{\sigma_t \in \Sigma} r(a_t'(\sigma_t), \sigma_t) \Pi(\sigma_t)$, equivalently for r_t'' , and so on. $\bar{r} := \sup_{a_t \in A} \sum_{\sigma_t \in \Sigma} r(a_t'(\sigma_t), \sigma_t) \Pi(\sigma_t)$ is an upper bound on r_t , whose existence follows from Assumption 4.

By the definition of dominance, there must exist an $\varepsilon > 0$ such that $W_t'' - W_t' \geq \varepsilon$ for all $t \geq s$. For all $t \geq s$, let $r_t''' \in [r_t', \bar{r}]$ be some number chosen so that the sequence $\{r_\tau'''\}_{\tau \geq t}$ satisfies the following inequality for all $t \geq s$:

$$\sum_{\tau=t} \beta^{\tau-t} (r_t''' - r_t') \in \left[\frac{\varepsilon}{2}, (W_t'' - W_t') \right)$$

$W_t'' - W_t' \geq \varepsilon$ implies that for all $t \geq s$ there must exist a $\tau \geq t$ such that $r_t' < \bar{r}$, so there is always scope to satisfy this inequality by a choice of ε sufficiently close to zero. By the normalisation in Section 3.4.1, it is always possible to find a sequence $\bar{\mathbf{a}}_s''' \in \mathcal{A}$ such that $\bar{\mathbf{a}}_s'''$ induces the promise sequence $\bar{\omega}_s''$ and implies a value for the policy criterion of r_t''' for all $t \geq s$. Now suppose the constraint set $\Xi^g(x'_{s-1})$ is expanded to $\Xi^g(x'_{s-1}) \cup (\bar{\mathbf{x}}_s', \bar{\mathbf{a}}_s''')$, and let W_t''' be the value of the discounted social welfare criterion in period t when $(\bar{\mathbf{x}}_s', \bar{\mathbf{a}}_s''')$ is implemented. By construction, $(\bar{\mathbf{x}}_t', \bar{\mathbf{a}}_t''')$ induces a promise sequence that is also satisfied by $(\bar{\mathbf{x}}_t'', \bar{\mathbf{a}}_t'')$ for all $t \geq s$, and $W_t''' < W_t''$ for all $t \geq s$, so $(\bar{\mathbf{x}}_t'', \bar{\mathbf{a}}_t'')$ constraint-dominates $(\bar{\mathbf{x}}_t', \bar{\mathbf{a}}_t''')$ for all $t \geq s$. Hence $\Xi^g(x'_{s-1}) \cup (\bar{\mathbf{x}}_s', \bar{\mathbf{a}}_s''')$ is an irrelevant extension of $\Xi^g(x'_{s-1})$. But $W_t''' - W_t' \geq \frac{\varepsilon}{2} > 0$ for all $t \geq s$, and $(\bar{\mathbf{x}}_s', \bar{\mathbf{a}}_s''')$ and $(\bar{\mathbf{x}}_s', \bar{\mathbf{a}}_s')$ imply the same state vector in every period, so $(\bar{\mathbf{x}}_s', \bar{\mathbf{a}}_s''') \succ^{TC} (\bar{\mathbf{x}}_s', \bar{\mathbf{a}}_s')$. Hence $(\bar{\mathbf{x}}_s', \bar{\mathbf{a}}_s')$ cannot belong to $D(x_{s-1})$.

Proof of Proposition 6

Let $\bar{\omega}_t'$ be the promise sequence induced by the Ramsey allocation, and consider the directional derivative $\delta_V(\bar{\omega}_t', x'_{t-1}; \bar{\mathbf{w}}_t)$ for some $t > s$. Rearranging the result in Proposition 13, this derivative will be given by:

$$\begin{aligned} \delta_V(\bar{\omega}_t', x'_{t-1}; \bar{\mathbf{w}}_t) &= - \int_{\sigma \in \Sigma} \lambda_t^k(\sigma) w_t(\sigma) d\Pi(\sigma) \\ &\quad + \sum_{\tau=t+1}^{\infty} \beta^{\tau-t} \int_{\sigma \in \Sigma} \{ [\lambda_{\tau-1}^m(\sigma) + \lambda_{\tau-1}^k(\sigma_-)] - \lambda_\tau^k(\sigma) \} w_\tau(\sigma) d\Pi(\sigma) \\ &= - \int_{\sigma \in \Sigma} \lambda_t^k(\sigma) w_t(\sigma) d\Pi(\sigma) \end{aligned}$$

where $w_t(\sigma) \in \mathbb{R}^J$ denotes the component of $\bar{\mathbf{w}}_s$ particular to date $t \geq s$ and state $\sigma \in \Sigma$, σ_- is the predecessor history to σ , and we have used the Ramsey optimality condition (11) to simplify. The result follows by noting that any vector of derivatives $\bar{\mathbf{w}}_s$ with $w_t(\sigma) < 0$ for all t and all σ in the specified positive-measure subset of Σ will deliver a marginal improvement in $V(\bar{\omega}_t', x'_{t-1})$ for all $t > s$, bounded above zero. Thus by Proposition 5,

$(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ cannot belong to $D(x'_{t-1})$ for any $t \geq s$.

Proof of Proposition 7

By Proposition 4, if $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ belongs to $D(x'_{t-1})$ for all $t \geq s$, $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ must solve Problem 1 for the promises that this allocation induces, denoted ω'_t . Thus by Proposition 5 it must be the case that ω'_t is undominated according to the ordering $\succ_{x'_{t-1}}^\omega$ for all $t \geq s$. Note also that the assumption V is differentiable at the chosen promise sequence implies that ω'_t must be strictly interior to $\Omega(x_{s-1})$. As above, let $\delta_V(\bar{\omega}'_s, x'_{s-1}; \bar{\mathbf{w}}_s)$ be the directional (Gateaux) derivative of $V(\bar{\omega}'_s, x'_{s-1})$ as $\bar{\omega}_s$ is varied along dimension $\bar{\mathbf{w}}_s$, and note that $\bar{\mathbf{w}}_s$ is required to be an element of the same vector space as $\bar{\omega}'_s$ (with $w_t(\sigma) \in \mathbb{R}^j$ denoting the component of $\bar{\mathbf{w}}_s$ particular to date $t \geq s$ and state $\sigma \in \Sigma$). If h is difference comparable then this is the space of promise sequences with bounded element-wise differences from one another. These differences will be invariant to any equivalent representation of the promises. If h is ratio comparable then the relevant space is the space of promise sequences with bounded ratio differences from one another. Again, these differences will be invariant to equivalent representations.

As Proposition 13 shows, the derivative at differentiable points can be written as:

$$\delta_V(\bar{\omega}'_s, x'_{s-1}; \bar{\mathbf{w}}_s) = \sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma \in \Sigma} \{ \beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] w_{t+1}(\sigma) - \lambda_t^k(\sigma) w_t(\sigma) \} d\Pi(\sigma)$$

Now fix some $\sigma \in \Sigma$, and suppose that there is a period τ such that the terms $[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]$ and $\lambda_t^k(\sigma)$ are both bounded above zero for all $t \geq \tau$. For each period t , consider the within-period component of the previous derivative expression, particular to σ :

$$\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] w_{t+1}(\sigma) - \lambda_t^k(\sigma) w_t(\sigma)$$

By the fact that the multiplier terms are bounded above zero, for any given $w_t(\sigma)$ it is possible to make the preceeding expression exceed any arbitrary constant $\varepsilon_t > 0$ by choosing $w_{t+1}(\sigma)$ to satisfy:

$$\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] w_{t+1}(\sigma) - \lambda_t^k(\sigma) w_t(\sigma) \geq \varepsilon_t \quad (48)$$

Difference-comparable h We first proceed under the assumption that h is difference comparable. In this case, the Gateaux derivative is defined for a bounded sequence $\{w_t(\sigma)\}_{t \geq \tau}$ for any $\sigma \in \Sigma$. If this sequence is such that inequality (48) can be satisfied for all $t \geq \tau$ for a sequence of ε_t values bounded above zero, and if this is true for all σ in a positive-measure subset of Σ , then the differential movement $\bar{\mathbf{w}}_\tau$ will generate a strict improvement for all policymakers from τ onwards, contradicting that ω'_t is undominated according to the ordering $\succ_{x'_{t-1}}^\omega$ for all $t \geq \tau$.

Suppose first that there is geometric convergence in the product $\prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$ to zero, i.e. for any $\tau \geq s$ there exists a $\rho \in (0, 1)$ and $K > 0$ such that for all $T > \tau$:

$$K\rho^{T-\tau} > \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$$

Then let $w_\tau(\sigma) > 0$, and for all $t \geq \tau$ set $w_{t+1}(\sigma) > 0$ recursively to satisfy the condition:

$$\frac{w_{t+1}(\sigma)}{w_t(\sigma)} \geq (1 + \gamma) \frac{\lambda_t^k(\sigma)}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \quad (49)$$

for some $\gamma > 0$ such that $\rho(1 + \gamma) < 1$, together with some lower bound $w_{t+1}(\sigma) \geq \underline{w} > 0$ and an upper bound $w_{t+1}(\sigma) \leq \bar{w} < \infty$. This upper bound is possible, because we have that:

$$K[\rho(1 + \gamma)]^{T-\tau} > (1 + \gamma)^{T-\tau} \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$$

and the object on the left-hand side converges to zero, whilst the existence of the lower bound is trivial. Given these values for the sequence $\{w_t(\sigma)\}_{t \geq \tau}$, set ε_t to satisfy:

$$\begin{aligned} & \frac{\lambda_t^k(\sigma)}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} + \frac{\varepsilon_t}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] w_t(\sigma)} \\ &= (1 + \gamma) \frac{\lambda_t^k(\sigma)}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \end{aligned}$$

or:

$$\varepsilon_t = \gamma \lambda_t^k(\sigma) w_t(\sigma)$$

Using this in (49) confirms that (48) is satisfied, and the bounds on $\lambda_t^k(\sigma)$ and $w_t(\sigma)$ imply ε_t is bounded above zero as required.

The alternative possibility when the multipliers are always strictly positive is that $K \left(\frac{1}{\rho}\right)^{T-\tau} < \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$ for some $K > 0$ and $\rho \in (0, 1)$. In this case choose some $\gamma \in (0, 1)$ sufficiently small that $\frac{(1-\gamma)}{\rho} > 1$. Let $w_\tau(\sigma) < 0$, and for all $t \geq \tau$ set $w_{t+1}(\sigma) < 0$ recursively so that the following is satisfied:

$$\frac{w_{t+1}(\sigma)}{w_t(\sigma)} \leq (1 - \gamma) \frac{\lambda_t^k(\sigma)}{\beta[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$$

together with the bound:

$$|w_{t+1}(\sigma)| \geq \underline{w}$$

for some $\underline{w} > 0$, and a similar upper bound. The existence of $\underline{w}$ follows from the fact

that:

$$0 < K \left(\frac{1-\gamma}{\rho} \right)^{T-\tau} < (1-\gamma)^{T-\tau} \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$$

for all T , and $\frac{1-\gamma}{\rho} > 1$. Now let ε_t be defined for all $t \geq \tau$ by:

$$\begin{aligned} & \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} + \frac{\varepsilon_t}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] w_t(\sigma)} \\ = & (1-\gamma) \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \end{aligned}$$

So that:

$$\varepsilon_t = -\gamma \lambda_t^k(\sigma) w_t(\sigma)$$

which is bounded above zero for all t . Thus there is a strict improvement in all periods, again contradicting that $\bar{\omega}'_t$ is undominated according to the ordering $\succ_{x'_{t-1}}^\omega$ for all $t \geq \tau$.

Ratio-comparable h When h is instead ratio comparable, the main formal adjustment to the proof is to take the Gateaux derivative as a bounded sequence of proportional deviations from the individual promises $\omega'_t(\sigma)$: $\{w_t(\sigma)\}_{t \geq \tau} = \{\tilde{w}_t(\sigma) \omega'_t(\sigma)\}_{t \geq \tau}$, with $\{\tilde{w}_t(\sigma)\}$ satisfying a uniform bound in t for any $\sigma \in \Sigma$. These proportional changes are independent of any admissible renormalisation by definition, and so can be generated by taking limits from alternative promise sequences that live in the same vector space as $\bar{\omega}'_\tau$.⁴⁰ Again, if this sequence of differential changes is such that inequality (48) can be satisfied for all $t \geq \tau$ for a sequence of ε_t values bounded above zero, and if this is true for all σ in a positive-measure subset of Σ , then ω'_t cannot be undominated for all $t \geq \tau$.

The argument then proceeds in a similar way to the difference comparable case. Suppose first that there is convergence in the product $\prod_{t=\tau}^{T-1} \left| \frac{\omega_t(\sigma)}{\omega_{t+1}(\sigma)} \right| \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$ to zero, i.e. for any $\tau \geq s$ there exists a $\rho \in (0, 1)$ and $K > 0$ such that for all $T > \tau$:

$$K \rho^{T-\tau} > \left| \frac{\omega_\tau(\sigma)}{\omega_T(\sigma)} \right| \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$$

Then choose an initial $\tilde{w}_\tau(\sigma)$ with $|\tilde{w}_\tau(\sigma)| > 0$ and $\text{sign}(\tilde{w}_\tau(\sigma)) = \text{sign}(\omega_\tau(\sigma))$, and for all $t \geq \tau$ set $\tilde{w}_{t+1}(\sigma)$ such that $\text{sign}(\tilde{w}_{t+1}(\sigma)) = \text{sign}(\omega_{t+1}(\sigma))$ and $\tilde{w}_{t+1}(\sigma)$ recursively satisfies:

$$\left| \frac{\tilde{w}_{t+1}(\sigma)}{\tilde{w}_t(\sigma)} \right| \geq (1+\gamma) \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \left| \frac{\omega_t(\sigma)}{\omega_{t+1}(\sigma)} \right| \quad (50)$$

for some $\gamma > 0$ such that $\rho(1+\gamma) < 1$, together with some lower bound $\tilde{w}_{t+1}(\sigma) \geq \underline{w} > 0$ and an upper bound $\tilde{w}_{t+1}(\sigma) \leq \tilde{w} < \infty$. This upper bound is possible, because we have

⁴⁰Part 1 of the Proposition implies $\omega_t(\sigma) > 0$, so the use of this as a reference point in defining the derivatives is not restrictive.

that:

$$K [\rho (1 + \gamma)]^{T-\tau} > (1 + \gamma)^{T-\tau} \left| \frac{\omega_\tau(\sigma)}{\omega_T(\sigma)} \right| \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$$

and the object on the left-hand side converges to zero, whilst the possibility of the lower bound is trivial. Note that if $\text{sign}(\omega_t(\sigma)) = \text{sign}(\omega_{t+1}(\sigma))$, condition (50) simply states:

$$\frac{\tilde{w}_{t+1}(\sigma)}{\tilde{w}_t(\sigma)} \geq (1 + \gamma) \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \frac{\omega_t(\sigma)}{\omega_{t+1}(\sigma)}$$

whereas if $\text{sign}(\omega_t(\sigma)) \neq \text{sign}(\omega_{t+1}(\sigma))$, it implies:

$$\frac{\tilde{w}_{t+1}(\sigma)}{\tilde{w}_t(\sigma)} \leq (1 + \gamma) \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \frac{\omega_t(\sigma)}{\omega_{t+1}(\sigma)}$$

Given the sequence $\{\tilde{w}_t(\sigma)\}_{t \geq \tau}$, set ε_t to satisfy:

$$\begin{aligned} & \frac{\lambda_t^k(\sigma) \omega_t(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \omega_{t+1}(\sigma)} + \frac{\varepsilon_t}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \omega_{t+1}(\sigma) \tilde{w}_t(\sigma)} \\ = & (1 + \gamma) \frac{\lambda_t^k(\sigma) \omega_t(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \omega_{t+1}(\sigma)} \end{aligned}$$

or:

$$\varepsilon_t = \gamma \lambda_t^k(\sigma) \omega_t(\sigma) \tilde{w}_t(\sigma)$$

Using this in (50), and multiplying through by $\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \omega_{t+1}(\sigma) \tilde{w}_t(\sigma)$, confirms that (48) is satisfied,⁴¹ and the bounds on $\lambda_t^k(\sigma) \omega_t(\sigma)$ and $\tilde{w}_t(\sigma)$ imply ε_t is bounded above zero as required.

The case where $K \left(\frac{1}{\rho}\right)^{T-\tau} < \left| \frac{\omega_\tau(\sigma)}{\omega_T(\sigma)} \right| \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$ for some $K > 0$ and $\rho \in (0, 1)$ can proceed by a symmetric adjustment to the proof from the difference-comparable case.

Proof of Proposition 8

The proof adopts the same approach as for Proposition 7, showing that a differential change to promises can generate an improvement in all periods when the stated conditions are not met. We first show that there must exist an $\alpha_t \geq 0$ such that the equality in the proof is almost always true, with the bounds on α_t established subsequently.

Suppose that there does not exist an α_t such that the equality in the proof is satisfied in t for Π -almost all σ . This means that in t there must be at least one degree of linear independence across (positive-measure values of) σ between the values of $[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]$ and of $\lambda_{t+1}^k(\sigma)$. Hence, under **difference comparability**, it is possible to find bounded

⁴¹Note that the sign of this expression will be negative if and only if $\text{sign}(\omega_t(\sigma)) \neq \text{sign}(\omega_{t+1}(\sigma))$.

differential changes $\{w_{t+1}(\sigma)\}_{\sigma \in \Sigma}$ such that the following two restrictions are met:

$$\int_{\sigma \in \Sigma} \lambda_{t+1}^k(\sigma) w_{t+1}(\sigma) d\Pi(\sigma) = 0 \quad (51)$$

$$\int_{\sigma \in \Sigma} [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] w_{t+1}(\sigma) d\Pi(\sigma) \geq 1 \quad (52)$$

Now, if the requirement in the proof is not satisfied, then it is possible to find a positive-measure subset of Σ that violates condition (14) by at least an amount ε at least every T periods, with $\varepsilon > 0$ and T finite. Thus at least every T periods it must be possible to satisfy conditions (51) and (52) with values for $w_t(\sigma)$ that are bounded in absolute value below some $\bar{w}$, uniform in t , and bounded away from zero for a positive-measure subset of Σ . Hence a strictly positive differential improvement is available of an amount at least equal to β^{T-1} in each period, applying the same logic as in the previous propositions. This contradicts that the original policy was time-consistently undominated, given Proposition 5. The case of **ratio comparability** proceeds on the same lines, normalising the derivatives by the promise values to preserve invariance.

We next show that $\alpha_t \leq 1$ can be imposed. Suppose otherwise. Then under **difference comparability** there must exist a $T < \infty$ such that for all σ in a positive-measure subset of Σ , in every period t there is a period $t + \tau$ with $\tau < T$ and:

$$\lambda_{t+\tau}^m(\sigma) + \lambda_{t+\tau}^k(\sigma_-) - \lambda_{t+\tau+1}^k(\sigma) \leq -\delta$$

for some $\delta > 0$. Now consider the differential change to $\bar{\omega}_s$ given by $\bar{\mathbf{w}}_s$ such that $w_{t+\tau}(\sigma) = -1$ for all date-states in which this inequality is true, and zero otherwise. In all period $t + \tau$ this delivers a differential improvement in σ -specific value given by:

$$-\lambda_{t+\tau}^m(\sigma) - \lambda_{t+\tau}^k(\sigma_-) + \lambda_{t+\tau+1}^k(\sigma) \geq \delta > 0$$

and in period $t + \tau + 1$ the improvement is:

$$\lambda_{t+\tau+1}^k(\sigma) \geq \delta > 0$$

Hence at any given t , for each state σ in the relevant subset of Σ there is a feasible differential improvement at least equal to $\beta^{T-1}\delta$. Since this is true in a positive-measure subset of Σ , the improvement is bounded above zero in value when assessed in any period $t \geq s$, so the original ω_s is dominated. A near-identical argument applies under **ratio comparability**, allowing for normalisation by the promise values.

Finally, we show that $\alpha_t > 0$ can be imposed. Then under **difference comparability** there must exist time periods t such that for a positive-measure subset of σ , the following

are true:

$$\begin{aligned}\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-) &\geq \delta \\ \lambda_{t+1}^k(\sigma) &= 0\end{aligned}$$

Consider a differential change $w_{t+1}(\sigma) = 1$, applied in all such date-states. The marginal value of this is at least $\beta\delta$ for state σ and date t . When the claim in the Proposition is not true, there is always a time period within T of the current date such that these gains can be realised for a positive-measure subset of states σ . Thus again there is a boundedly-positive marginal improvement available in net present value at every point in time, contradicting that the original policy was time-consistently undominated. Again, the case of **ratio comparability** proceeds symmetrically, setting $w_{t+1}(\sigma) = |\omega_{t+1}(\sigma)|$ in this case, to preserve invariance.

Proof of Proposition 9

We present the main proof under **difference comparability**. Quasiconcavity of the value function implies, by the usual logic, that the absence of marginal gains from moving allocations along a given vector dimension will also ensure the absence of discrete gains. Thus, applying Proposition 5, it is sufficient to show that when the three specified conditions are satisfied, there is no marginal change to the promises $\bar{\mathbf{w}}_s$ such that $\delta(V(\bar{\omega}'_t, x'_{t-1}), \bar{\mathbf{w}}_t)$ will be bounded above zero for all t sufficiently large, including at the limit as $t \rightarrow \infty$.

We start with two definitions. It aids the proof to define the scalar η_t for $t \geq s$ recursively by:

$$\eta_s : = 1$$

and for $t > s$:

$$\eta_t = \frac{\alpha_{t-1}}{\beta} \eta_{t-1}$$

Note that $\eta_t > 0$ for all t , since $\alpha_t \in (0, 1)$.

In addition, for all $t \geq s$ and $\sigma \in \Sigma$, define $\Delta_t(\sigma)$ as a measure of the deviation from the limit in Condition 2(b):

$$\frac{\lambda_{t+1}^k(\sigma)(1 + \Delta_t(\sigma))}{\alpha_t [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} \equiv 1 \tag{53}$$

Note that linear convergence implies that the product:

$$\prod_{\tau=t}^r (1 + \Delta_{\tau}(\sigma))$$

converges to a finite positive constant as $r \rightarrow \infty$.

Condition 2(a) in the Proposition states that the following object is bounded above zero and below ∞ , uniformly in r , for all $\sigma \in \Sigma$:

$$\prod_{t=\tau}^{r-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]}$$

Applying the identity (53), this can be rewritten as follows:

$$\begin{aligned} \prod_{t=\tau}^{r-1} \frac{\lambda_t^k(\sigma)}{\beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)]} &= \prod_{t=\tau}^{r-1} \frac{\alpha_t}{\beta} \frac{\lambda_t^k(\sigma)}{\lambda_{t+1}^k(\sigma) (1 + \Delta_t(\sigma))} \\ &= \frac{\lambda_{\tau}^k(\sigma)}{\lambda_r^k(\sigma)} \prod_{t=\tau}^{r-1} \frac{\eta_{t+1}}{\eta_t (1 + \Delta_t(\sigma))} \\ &= \frac{\lambda_{\tau}^k(\sigma)}{\lambda_r^k(\sigma)} \frac{\eta_r}{\eta_{\tau}} \prod_{t=\tau}^{r-1} \frac{1}{(1 + \Delta_t(\sigma))} \end{aligned} \quad (54)$$

Since the final product term in $(1 + \Delta_t(\sigma))$ is bounded, it follows that the object $\frac{\lambda_{\tau}^k(\sigma)}{\lambda_r^k(\sigma)} \frac{\eta_r}{\eta_{\tau}}$ must likewise be bounded above zero and below ∞ , uniformly in r , for all τ

We can further define $\tilde{\lambda}_t^k(\sigma)$ by:

$$\tilde{\lambda}_t^k(\sigma) := \frac{1}{\eta_t} \lambda_t^k(\sigma)$$

Notice that the boundedness of $\frac{\lambda_{\tau}^k(\sigma)}{\lambda_r^k(\sigma)} \frac{\eta_r}{\eta_{\tau}}$ implies that $\tilde{\lambda}_t^k(\sigma)$ is bounded above zero and below ∞ in t for all σ , irrespective of the convergence properties of $\lambda_t^k(\sigma)$ and η_t .

Now suppose, contrary to the claim in the Proposition, that there exists an alternative promise sequence $\bar{\omega}_s$ that is bounded away from the lower contour set of $\bar{\omega}'_s$ for all $t \geq \tau$, and some $\tau \geq s$. We first translate this into a derivative statement. As shown in Proposition 13, the Gateaux derivative in all periods $t \geq \tau$ satisfies:

$$\delta_V(\bar{\omega}'_t, x'_{t-1}; \bar{\mathbf{w}}_t) = \sum_{r=t}^{\infty} \beta^{r-t} \int_{\sigma \in \Sigma} \{ \beta [\lambda_r^m(\sigma) + \lambda_r^k(\sigma_-)] w_{r+1}(\sigma) - \lambda_r^k(\sigma) w_r(\sigma) \} d\Pi(\sigma)$$

where $w_r(\sigma)$ is the marginal increase in the date-state-specific promise $\omega'_r(\sigma)$. Using the

definition of $\Delta_r(\sigma)$ above, this can be rewritten as:

$$\delta_V(\bar{\omega}'_t, x'_{t-1}; \bar{\mathbf{w}}_t) = \sum_{r=t}^{\infty} \beta^{r-t} \int_{\sigma \in \Sigma} \left\{ \beta \lambda_{r+1}^k(\sigma) \frac{1}{\alpha_r} [1 + \Delta_r(\sigma)] w_{r+1}(\sigma) - \lambda_r^k(\sigma) w_r(\sigma) \right\} d\Pi(\sigma)$$

Since $\lambda_t^k(\sigma)$ is not guaranteed to be bounded (above or below) in t , this object need not be bounded, which in general will not allow us to reason from finite differential improvements in V to boundedly positive gains in promise space and vice-versa.⁴² To overcome this, we can normalise it by η_t , giving:

$$\frac{\delta_V(\bar{\omega}'_t, x'_{t-1}; \bar{\mathbf{w}}_t)}{\eta_t} = \sum_{r=t}^{\infty} \beta^{r-t} \frac{\eta_r}{\eta_t} \int_{\sigma \in \Sigma} \left\{ \tilde{\lambda}_{r+1}^k(\sigma) [1 + \Delta_r(\sigma)] w_{r+1}(\sigma) - \tilde{\lambda}_r^k(\sigma) w_r(\sigma) \right\} d\Pi(\sigma)$$

This is equivalent to rescaling the value function in t by the factor $\frac{1}{\eta_t}$. By the recursive definition of η_t , we have:

$$\beta^{r-t} \frac{\eta_r}{\eta_t} = \prod_{\tau=t}^{r-1} \alpha_{\tau} < \bar{\alpha}^{r-t}$$

with $\bar{\alpha} < 1$, so the boundedness of $\tilde{\lambda}_{r+1}^k(\sigma)$ and convergence of $\Delta_r(\sigma)$ (Condition 2(b)) implies we have $\lim_{t \rightarrow \infty} \sup_{\bar{\mathbf{w}}_t} \left[\frac{\delta_V(\bar{\omega}'_t, x'_{t-1}; \bar{\mathbf{w}}_t)}{\eta_t} \right] < \infty$, recalling that $w_t(\sigma)$ must be bounded uniformly in t by the definition of the derivative. It follows from the concavity of V that the alternative promise sequence $\bar{\omega}_t'' := (\bar{\omega}'_t + \alpha \bar{\mathbf{w}}_t)$ is bounded away from the lower contour set of $\bar{\omega}_t'$ for all $t \geq \tau$ only if $\frac{\delta_V(\bar{\omega}'_t, x'_{t-1}; \bar{\mathbf{w}}_t)}{\eta_t} \geq \varepsilon$ holds for all $t \geq \tau$, some $\tau \geq s$ and $\varepsilon > 0$. Thus an improvement requires a derivative vector with the property:

$$\sum_{r=t}^{\infty} \beta^{r-t} \frac{\eta_r}{\eta_t} \int_{\sigma \in \Sigma} \left\{ \tilde{\lambda}_{r+1}^k(\sigma) [1 + \Delta_r(\sigma)] w_{r+1}(\sigma) - \tilde{\lambda}_r^k(\sigma) w_r(\sigma) \right\} d\Pi(\sigma) \geq \varepsilon \quad (55)$$

for all t sufficiently large. Note that by Condition 2(b), for every $\epsilon > 0$ there is r sufficiently large that $|\Delta_r(\sigma)| < \epsilon$ for all $\sigma \in \Sigma$.

Rewriting (55) gives:

$$\begin{aligned} & \sum_{r=t+1}^{\infty} \beta^{r-(t+1)} \left(\frac{\beta}{\alpha_{r-1}} - \beta \right) \frac{\eta_r}{\eta_t} \int_{\sigma \in \Sigma} \tilde{\lambda}_r^k(\sigma) w_r(\sigma) d\Pi(\sigma) \\ & \geq \int_{\sigma \in \Sigma} \tilde{\lambda}_t^k(\sigma) w_t(\sigma) d\Pi(\sigma) + \varepsilon \\ & \quad - \sum_{r=t}^{\infty} \beta^{r-t} \int_{\sigma \in \Sigma} \frac{\eta_r}{\eta_t} \Delta_r(\sigma) \tilde{\lambda}_{r+1}^k(\sigma) w_{r+1}(\sigma) d\Pi(\sigma) \end{aligned} \quad (56)$$

⁴²If $\lambda_t^k(\sigma) \rightarrow \infty$ it is possible that a boundedly large increase in V could be achieved by a change in promise values away from $\bar{\omega}'_t$ that becomes vanishingly small as $t \rightarrow \infty$. This would not satisfy the requirement for the improving promise sequence to be bounded away from the previous sequence at the limit as $t \rightarrow \infty$.

Using the definition of η_t , the last term here simplifies to:

$$\sum_{r=t}^{\infty} \left[\prod_{\tau=t}^{r-1} \alpha_{\tau} \right] \int_{\sigma \in \Sigma} \tilde{\lambda}_{r+1}^k(\sigma) w_{r+1}(\sigma) \Delta_r(\sigma) d\Pi(\sigma)$$

Since $\alpha_t \leq \bar{\alpha} < 1$ and $\tilde{\lambda}_t^k(\sigma)$ and $w_t(\sigma)$ are both bounded uniformly in t , this expression converges to 0 as $\Delta_t(\sigma)$ does so across σ . Thus it is possible to find a sufficiently large T such that:

$$\left| \sum_{r=t}^{\infty} \left[\prod_{\tau=t}^{r-1} \alpha_{\tau} \right] \int_{\sigma \in \Sigma} \tilde{\lambda}_{r+1}^k(\sigma) w_{r+1}(\sigma) \Delta_r(\sigma) d\Pi(\sigma) \right| < \frac{\varepsilon}{2}$$

for all $t \geq T$. Using this in inequality (56) implies that for sufficiently large t we have:

$$\begin{aligned} & \sum_{r=t+1}^{\infty} \beta^{r-(t+1)} \left(\frac{\beta}{\alpha_{r-1}} - \beta \right) \frac{\eta_r}{\eta_t} \int_{\sigma \in \Sigma} \tilde{\lambda}_r^k(\sigma) w_r(\sigma) d\Pi(\sigma) \\ & \geq \int_{\sigma \in \Sigma} \tilde{\lambda}_t^k(\sigma) w_t(\sigma) d\Pi(\sigma) + \frac{\varepsilon}{2} \end{aligned}$$

Now consider the sum:

$$\sum_{r=t+1}^{\infty} \beta^{r-(t+1)} \left(\frac{\beta}{\alpha_{r-1}} - \beta \right) \frac{\eta_r}{\eta_t}$$

Since $\alpha_r \in (0, 1)$ and $\eta_r > 0$ for all r , each element of this sum is positive. In addition, we have:

$$\begin{aligned} & \sum_{r=t+1}^{\infty} \beta^{r-(t+1)} \left(\frac{\beta}{\alpha_{r-1}} - \beta \right) \frac{\eta_r}{\eta_t} \\ & = \sum_{r=t+1}^{\infty} \beta^{r-(t+1)} \left(\frac{\eta_{r-1}}{\eta_t} - \beta \frac{\eta_r}{\eta_t} \right) \\ & = \sum_{r=t}^{\infty} \beta^{r-t} \frac{\eta_r}{\eta_t} - \sum_{r=t+1}^{\infty} \beta^{r-t} \frac{\eta_r}{\eta_t} \\ & = \frac{\eta_t}{\eta_t} = 1 \end{aligned}$$

Thus the sum can be interpreted as a probability distribution weighting time periods, and the expression:

$$\sum_{r=t+1}^{\infty} \beta^{r-(t+1)} \left(\frac{\beta}{\alpha_{r-1}} - \beta \right) \frac{\eta_r}{\eta_t} \int_{\sigma \in \Sigma} \tilde{\lambda}_r^k(\sigma) w_r(\sigma) d\Pi(\sigma)$$

is a weighted average of values for $\int_{\sigma \in \Sigma} \tilde{\lambda}_r^k(\sigma) w_r(\sigma) d\Pi(\sigma)$ across periods $r > t$. Inequality (56) states that this weighted average always exceeds the value of the same object in

t itself, by at least an amount $\frac{\varepsilon}{2} > 0$. This is possible only if the object:

$$\int_{\sigma \in \Sigma} \tilde{\lambda}_t^k(\sigma) w_t(\sigma) d\Pi(\sigma)$$

is growing without bound in t . But this is inconsistent with boundedness of $w_t(\sigma)$ and $\tilde{\lambda}_t^k(\sigma)$. The former of these is a necessary requirement for the improving promise sequence $\bar{\omega}_t''$ to be well defined in the chosen vector space, and the latter was established above. Hence we have a contradiction.

The proof under **ratio comparability** proceeds near-identically, allowing for the fact that Gateaux derivatives can now only be established as bounded ratio changes in promises: $\{w_t(\sigma)\}_{t \geq \tau} = \{\omega_t(\sigma) \tilde{w}_t(\sigma)\}_{t \geq \tau}$, with $\{\tilde{w}_t(\sigma)\}_{t \geq \tau}$ satisfying a uniform bound.

B Properties of the value function

The implication of Proposition 5 is that undominated allocations under the ordering $\succeq^{TC}$ can be identified by reference to preferences across promise sequences, $\succeq_{x_{s-1}}^\omega$. These preferences are defined on the space of promise sequences for which Problem 1 has a solution. They can be represented by the value function associated with Problem 1. This representation provides an important step in operationalising our approach. This Appendix analyses the properties of its two components: the feasible set of promises, and the value function.

B.1 Feasible promise sequences

For some choices of $\bar{\omega}_s$ the constraint set for Problem 1 may be empty – there simply does not exist an allocation that can make good on these promises. Clearly these are not feasible selections. The set of feasible promise sequences from s onwards is denoted by $\Omega(x_{s-1})$:

$$\Omega(x_{s-1}) := \{\bar{\omega}_s \in \mathcal{W} : \text{constraint set to Problem 1 nonempty, given } x_{s-1} \in X\}$$

To analyse differential changes to promises, it is useful to restrict attention to the interior of $\Omega(x_{s-1})$. This is denoted by $\overset{\circ}{\Omega}(x_{s-1})$.

Convexity of $\Omega(x_{s-1})$ An important regularity property to be able to place on $\Omega(x_{s-1})$ is convexity. The next Proposition establishes the conditions under which this will hold.

Proposition 11. *Suppose Assumptions 4 and 5 hold. For any $x_s \in X$, the space $\Omega(x_{s-1})$ is **convex**.*

The proof of this is omitted to avoid repetition: the result follows directly from arguments contained in the more general proof of Proposition 12 below. Note that the concavity of r will not be needed for this result.

B.2 Value of the inner problem

The maximised **value** of the inner problem is denoted by $V(\bar{\omega}_s; x_{s-1})$, for all $\bar{\omega}_s \in \Omega(x_{s-1})$ and all $x_{s-1} \in X$. For all $\omega_s \in \mathcal{W}$ not in $\Omega(x_{s-1})$, we normalise $V(\bar{\omega}_s; x_{s-1})$ to $-\infty$ for convenience. Note that V can be viewed as a cardinalisation of the preference ordering $\succeq_{x_{s-1}}^\omega$, given x_{s-1} . Thus time-consistently undominated promise choices can be investigated by reference to the effect of promises on the value of V at every horizon.

Concavity of V As in conventional optimisation theory, a particularly useful property for V to exhibit is concavity, as this allows the application of global methods for constrained optimisation. The next Proposition establishes the conditions under which concavity will hold.

Proposition 12. *Suppose Assumptions 4 and 5 hold. For all $x_{s-1} \in X$, $V(\cdot; x_{s-1})$ is **concave** in $\bar{\omega}_s \in \Omega(x_{s-1})$.*

Proof. To ease notation we suppress the dependence of functions on σ . Consider two promise sequences $\bar{\omega}'_s, \bar{\omega}''_s \in \Omega(x_{s-1})$. To establish concavity we must show:

$$V(\alpha \bar{\omega}'_s + (1 - \alpha) \bar{\omega}''_s; x_{s-1}) \geq \alpha V(\bar{\omega}'_s; x_{s-1}) + (1 - \alpha) V(\bar{\omega}''_s; x_{s-1}) \quad (57)$$

for all $\alpha \in (0, 1)$. Let $\bar{\mathbf{y}}' := (\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ and $\bar{\mathbf{y}}'' := (\bar{\mathbf{x}}''_s, \bar{\mathbf{a}}''_s)$ solve Problem 1 for $\bar{\omega}'_s$ and $\bar{\omega}''_s$ respectively. It follows from the concavity of r (Assumption 5) that (57) must be satisfied provided the convex combination $\alpha \bar{\mathbf{y}}' + (1 - \alpha) \bar{\mathbf{y}}''$ is feasible when the promise sequence is $\alpha \bar{\omega}'_s + (1 - \alpha) \bar{\omega}''_s$. In this case the feasible selection $\alpha \bar{\mathbf{y}}' + (1 - \alpha) \bar{\mathbf{y}}''$ will deliver a value at least as great as the right-hand side of (57), which is then a lower bound on $V(\alpha \bar{\omega}'_s + (1 - \alpha) \bar{\omega}''_s; x_{s-1})$. The quasiconcavity of g implies that if (4) is satisfied in all time periods by both $\bar{\mathbf{y}}'$ and $\bar{\mathbf{y}}''$ then it must also be satisfied by $\alpha \bar{\mathbf{y}}' + (1 - \alpha) \bar{\mathbf{y}}''$. These constraints are unaffected by variations in the promise values. Thus it remains only to show that constraints (8) and (9) are also satisfied by the convex combination. Consider (8). For all $t \geq s$, we need:

$$h(\alpha a'_t + (1 - \alpha) a''_t) + \beta [\alpha \omega'_{t+1} + (1 - \alpha) \omega''_{t+1}] \geq h^0(\alpha a'_t + (1 - \alpha) a''_t)$$

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

$$\alpha h(a'_t) + (1 - \alpha) h(a''_t) + \beta [\alpha \omega'_{t+1} + (1 - \alpha) \omega''_{t+1}] \geq \alpha h^0(a'_t) + (1 - \alpha) h^0(a''_t)$$

But by concavity of h :

$$h(\alpha a'_t + (1 - \alpha) a''_t) \geq \alpha h(a'_t) + (1 - \alpha) h(a''_t)$$

and by convexity of h^0 :

$$h^0(\alpha a'_t + (1 - \alpha) a''_t) \leq \alpha h^0(a'_t) + (1 - \alpha) h^0(a''_t)$$

Collecting together, this establishes the desired inequality. An identical argument confirms that (9) is likewise satisfied for all $t \geq s$. Thus $\alpha \bar{\mathbf{y}}' + (1 - \alpha) \bar{\mathbf{y}}''$ is feasible when the promise sequence is $\alpha \omega'_s + (1 - \alpha) \omega''_s$, completing the proof. $\square$

Placing this additional structure on V does not come without a cost. Assumption 5

requires concavity in the h function and convexity in the h^0 function – not just quasiconcavity/convexity. In many problems of interest these will not be easy to guarantee. As ever, when the required assumptions are not satisfied the analysis can proceed, but with caveats. The most direct analogy in this case is with the analysis of consumer demand when the utility function is not known to be quasi-concave.

Concave, real-valued functions of a real interval are well known to have appealing continuity properties. The following corollary is a standard result:

Corollary 3. *Suppose the assumptions for Proposition 12 are true. Fix $x_{s-1} \in X$, and let $\bar{\omega}'_s$ and $\bar{\omega}''_s$ be arbitrary selections from $\Omega(x_{s-1})$. Then $V(\alpha\bar{\omega}'_s + (1-\alpha)\bar{\omega}''_s; x_{s-1})$ is continuous in $\alpha \in [0, 1]$, has left derivatives with respect to α for all $\alpha \in (0, 1]$, and has right derivatives with respect to α for all $\alpha \in [0, 1)$. These derivatives coincide for almost all $\alpha \in (0, 1)$.*

This provides a solid basis for taking directional derivatives of V with respect to the promise sequence.

Derivatives of V The analysis that follows will characterise time-consistently undominated policy by reference to the slope of the V function as promises are varied. For $x_{s-1} \in X$ and $\bar{\omega}_s \in \Omega(x_{s-1})$, the directional (Gateaux) derivative of V is denoted by $\delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s)$, defined for all $\bar{\mathbf{w}}_s \in \mathcal{W}$ by:⁴³

$$\delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s) := \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} [V(\bar{\omega}'_s + \alpha\bar{\mathbf{w}}_s; x'_{s-1}) - V(\bar{\omega}'_s; x'_{s-1})]$$

wherever this limit exists. Where V is not differentiable in the relevant dimension, $\delta_V^+(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s)$ will denote the above limit as $\alpha \rightarrow 0$ from above, and $\delta_V^-(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s)$ as $\alpha \rightarrow 0$ from below.

Where V is differentiable, the usual envelope results for value functions will apply, so that the derivatives of V will be defined in terms of Lagrange multipliers on the promise-keeping and promise-making constraints.

In general we denote the present-value multiplier on promise-keeping constraint (9) for history σ in period t by $\lambda_t^k(\sigma)$, and the corresponding promise-making constraint (8) by $\lambda_t^m(\sigma)$. Consistent with earlier notation, λ_t^k and λ_t^m are the collection of within-period multipliers across $\sigma \in \Sigma$, and $\bar{\lambda}_s^k$ and $\bar{\lambda}_s^m$ are infinite sequences of these from s on. The space that $\bar{\lambda}_s^k$ and $\bar{\lambda}_s^m$ inhabit is denoted $\mathcal{W}^*$.

Confirming the existence of Lagrange multipliers in convex optimisation problems generally requires the existence of a point that is strictly interior to the constraint set.⁴⁴ Formally, we will make use of the following:

⁴³The individual component of $\bar{\mathbf{w}}_s$ for period t and state σ is denoted $w_t(\sigma)$ in what follows, consistent with the notation for promises.

⁴⁴See, for instance, Luenberger (1969), §8.3, Theorem 1.

Definition. For any $x_{s-1} \in X$ and $\bar{\omega}_s \in \Omega(x_{s-1})$, we say that the corresponding constraint set for Problem 1 contains an **interior point** if there is an allocation $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ in this constraint set that satisfies the following two inequalities for Π -almost all $\sigma \in \Sigma$ and all $t \geq s$:

$$\begin{aligned}\mathbb{E}_{t-1}[h(a'(\sigma'), \sigma') + \beta \omega_{t+1}(\sigma') | \sigma] - \omega_t(\sigma) &\geq \varepsilon \\ h(a'_t(\sigma), \sigma) + \beta \omega_{t+1}(\sigma) - h^0(a'_t(\sigma), \sigma) &\geq \varepsilon\end{aligned}$$

for some $\varepsilon > 0$, independent of σ and t .

The existence of an inner point is not a trivial requirement. It is immediate, for instance, that it cannot be satisfied when $\bar{\omega}_s$ lies at the boundary of $\Omega(x_{s-1})$.⁴⁵ In addition, the condition rules out the simple incorporation of equality constraints as two-sided inequalities, since in this case interiority is impossible. Extensions to the main arguments are possible that allow for linear forward-looking constraints, but we neglect these to avoid over-complicating the analysis.⁴⁶

Proposition 13. *Suppose Assumptions 1, 4 and 5 hold. Fix $x_{s-1} \in X$, and let $\bar{\omega}_s \in \Omega(x_{s-1})$ be such that the constraint set for Problem 1 contains an interior point. Then wherever the directional derivative $\delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s)$ exists and is finite-valued, there is a pair of Lagrange multiplier sequences $\bar{\lambda}_s^{\mathbf{k}}$ and $\bar{\lambda}_s^{\mathbf{m}}$ in $\mathcal{W}^*$ such that $\delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s)$ is given by:*

$$\delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s) = \sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma \in \Sigma} \{ \beta [\lambda_t^{\mathbf{m}}(\sigma) + \lambda_t^{\mathbf{k}}(\sigma_-)] w_{t+1}(\sigma) - \lambda_t^{\mathbf{k}}(\sigma) w_t(\sigma) \} d\Pi(\sigma) \quad (58)$$

with σ_- the predecessor history to σ .

Proof. Existence of the saddle point multipliers follows from direct application of Theorem 1, §8.3 in Luenberger (1969), given assumptions 4 and 5. The promise-keeping constraint can be rewritten for all t and σ as:

$$\mathbb{E}_{t-1}[h(a_t(\sigma'), \sigma') | \sigma] \geq \gamma(\sigma)$$

where $\gamma_t(\sigma) := \omega_t(\sigma) - \beta \mathbb{E}_t[\omega_{t+1}(\sigma') | \sigma]$, so that the vector movement in promises $\bar{\mathbf{w}}_s$ causes a per-unit change in $\gamma_t(\sigma)$ of $w_t(\sigma) - \beta \mathbb{E}_t[w_{t+1}(\sigma') | \sigma]$. Hence, applying Theorem

⁴⁵If it were, then a sufficiently small change in promises in any direction would be consistent with the existence of a feasible allocation. Hence we could not be at the boundary.

⁴⁶Linear forward-looking equality constraints will most commonly arise in linear-quadratic problems, in which case conventional techniques from linear analysis can provide an equivalent characterisation.

1, §8.5 of Luenberger (1969), where the derivative exists it is given by:

$$\begin{aligned}
\delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s) &= \sum_{t=s}^{\infty} \beta^{t-s} \left\{ \int_{\sigma \in \Sigma} \beta \lambda_t^m(\sigma) w_{t+1}(\sigma) d\Pi(\sigma) \right. \\
&\quad \left. + \int_{\sigma \in \Sigma} \lambda_t^k(\sigma) [\beta \mathbb{E}_t[w_{t+1}(\sigma') | \sigma] - w_t(\sigma)] d\Pi(\sigma) \right\} \\
&= \sum_{t=s}^{\infty} \beta^{t-s} \left\{ \int_{\sigma \in \Sigma} [\beta \lambda_t^m(\sigma) w_{t+1}(\sigma) - \lambda_t^k(\sigma) w_t(\sigma)] d\Pi(\sigma) \right. \\
&\quad \left. + \beta \int_{\sigma \in \Sigma} \int_{\sigma' \in \Sigma} \lambda_t^k(\sigma) w_{t+1}(\sigma') d\Pi(\sigma' | \sigma) d\Pi(\sigma) \right\} \\
&= \sum_{t=s}^{\infty} \beta^{t-s} \left\{ \int_{\sigma \in \Sigma} [\beta \lambda_t^m(\sigma) w_{t+1}(\sigma) - \lambda_t^k(\sigma) w_t(\sigma)] d\Pi(\sigma) \right. \\
&\quad \left. + \beta \int_{\sigma \in \Sigma} \lambda_t^k(\sigma_-) w_{t+1}(\sigma) d\Pi(\sigma) \right\}
\end{aligned}$$

where the last line applies Assumption 1, and σ_- is the predecessor history to σ . This delivers the stated expression under differentiability. (The right and left derivatives without differentiability, discussed in the text, follow from identical logic, combined with the concavity of V .) $\square$

All of the major characterisation results that follow will assume differentiability in V , applying condition (58), but a generalisation to points of non-differentiability would be technically straightforward. Where the derivative $\delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s)$ does not exist, there is a set of Lagrange multipliers $\mathbf{\Lambda}_s^k \times \mathbf{\Lambda}_s^m \subset \mathcal{W}^* \times \mathcal{W}^*$ such that $\delta_V^+(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s)$ is the minimum in $\mathbf{\Lambda}_s^k \times \mathbf{\Lambda}_s^m$ of the object on the right-hand side of (58), and $\delta_V^-(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s)$ is its maximum.

C Problems with incentive compatibility constraints

In this section we sketch the arguments needed to extend the main characterisation results, Propositions 7 to 9, to problems with incentive compatibility restrictions of the form (6). This extension is used in the example of Section 9.3 in the main text, applying the model of Atkeson and Lucas (1992). In all cases the extension is quite mechanical, but expanded dimensionality in the constraint set makes the notation more burdensome. This is why we relegate the treatment to an appendix.

Constraint (6) states:

$$\mathbb{E}_t \left[\sum_{\tau=0}^{\infty} \beta^{\tau} h(a_{t+\tau}(\sigma_{t+\tau}), \sigma_{t+\tau}) \middle| \sigma_t \right] \geq \mathbb{E}_t \left[\sum_{\tau=0}^{\infty} \beta^{\tau} h(a_{t+\tau}(\tilde{\sigma}_{t+\tau}), \sigma_{t+\tau}) \middle| \sigma_t \right] \quad (59)$$

We assume that σ_t is an infinite history of draws of some stochastic variable $\theta \in \Theta \subseteq \mathbb{R}$. These draws are taken to be iid through time, so σ_t is not informative about the expected sequence of draws from $t+1$ on. Since incentive compatibility relates to the period-by-period reporting decision, it is helpful to represent σ_t as (σ_{t-1}, θ_t) , where σ_{t-1} is an infinite past history of θ draws, and θ_t is the current realisation. The iid assumption means that it does not matter in t whether σ_{t-1} was a true history or simply a reported one. We denote by $\pi(\theta)$ the time-invariant density function of θ .

Given this, constraint (6) can be decomposed into promise-making and promise-keeping restrictions as follows, for all $(\sigma_-, \theta) \in \Sigma$:

$$h(a_t(\sigma_-, \theta), \theta) + \beta \omega_{t+1}(\sigma_-, \theta) \geq h(a_t(\sigma_-, \tilde{\theta}), \theta) + \beta \omega_{t+1}(\sigma_-, \tilde{\theta}) \quad (60)$$

$$\mathbb{E}_{t-1}[h(a_t(\sigma_-, \theta), \theta) + \beta \omega_{t+1}(\sigma_-, \theta) | \sigma_-] = \omega_t(\sigma_-) \quad (61)$$

where (60) must hold for all $\tilde{\theta} \in \Theta$.⁴⁷ Note that the promise-keeping constraint (61) must now be stated with equality: providing utility in excess of the required value may violate incentive compatibility for those with alternative histories. These constraints and the feasibility restrictions in (4) define an equivalent to the inner problem, Problem 1. The value function associated with this problem can be denoted $V(\bar{\omega}_s, x_{s-1})$ as before. The multiplier on constraint (60) is denoted $\lambda_t^m(\sigma_-, \theta, \tilde{\theta})$, and on (61) it is $\lambda_t^k(\sigma_-)$, normalised by the relative measure of the σ and θ draws in both cases. Proceeding as before, it is easy to show that if V is differentiable in the promise sequence then its

⁴⁷In general only a small subset of these constraints will be binding at a chosen allocation.

directional derivative along the dimension $\bar{\mathbf{w}}_s$ is given by:

$$\begin{aligned} \delta_V(\bar{\omega}_s, x_{s-1}; \bar{\mathbf{w}}_s) = & \sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma \in \Sigma} \left\{ \beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) \right. \right. \\ & \left. \left. - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right] w_{t+1}(\sigma) \right. \\ & \left. - \lambda_t^k(\sigma) w_t(\sigma) \right\} d\Pi(\sigma) \end{aligned}$$

where $\sigma := (\sigma_-, \theta)$. There are two main extensions here relative to the case in the main text. First, a marginal increase in the promise value $\omega_{t+1}(\sigma_-, \theta)$ in principle relaxes an entire *set* of promise-making constraints that are of relevance to agents whose true draw is (σ_-, θ) – hence the first integral across alternative $\tilde{\theta}$ reports. Second, an increase in this promise value tightens the set of promise-making constraints for agents with the common past (reported) history σ_- , across all current $\tilde{\theta}$ draws. The second of these ensures that the term in square brackets in the derivative expression need not be positive, which necessitates some adjustments to the analysis. Up to this qualification, characterisation results will proceed as before.

The equivalent statement to Proposition 7 now is:

Proposition 14. *Suppose that the policy $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ is time-consistently undominated, given some initial $x'_{s-1} \in X$, and assume that V is differentiable at the induced promise sequence $\bar{\omega}'_s$. When h is difference-comparable,⁴⁸ for Π -almost all $\sigma \in \Sigma$, either:*

1. *There is no period τ such that both $\left| \int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right|$ and $|\lambda_t^k(\sigma)|$ are bounded above zero for all $t \geq \tau$.*

or :

2. *For all $\rho \in (0, 1)$ and all positive scalars K_1 and K_2 , it is possible to find a $\tau \geq s$ and $T > \tau$ such that:*

$$K_1 \rho^{r-\tau} < \prod_{t=\tau}^{r-1} \left| \frac{\lambda_t^k(\sigma)}{\beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right]} \right| < K_2 \left(\frac{1}{\rho} \right)^{r-\tau}$$

for all $r \geq T$.

Proof. The proof mimics the difference comparable case above, with minor adjustments.

⁴⁸Incentive compatibility constraints are commonly based on difference-comparable dynamic utility functions, so we omit the case of ratio comparability.

Suppose that there is geometric convergence in the product:

$$\prod_{t=\tau}^{T-1} \left| \frac{\lambda_t^k(\sigma)}{\beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right]} \right|$$

to zero, i.e. for any $\tau \geq s$ there exists a $\rho \in (0, 1)$ and $K > 0$ such that for all $T > \tau$:

$$K\rho^{T-\tau} > \prod_{t=\tau}^{T-1} \left| \frac{\lambda_t^k(\sigma)}{\beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right]} \right|$$

Let τ be such that both $\left| \int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right|$ and $|\lambda_t^k(\sigma)|$ are bounded above zero for all $t \geq \tau$ – i.e., part 1 of the Proposition is not true. Then choose $w_\tau(\sigma) > 0$ arbitrarily, and for all $t \geq \tau$ set $w_{t+1}(\sigma)$ recursively to satisfy the condition:

$$\beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right] w_{t+1}(\sigma) \geq (1 + \gamma) |\lambda_t^k(\sigma) w_t(\sigma)| \quad (62)$$

for some $\gamma > 0$ such that $\rho(1 + \gamma) < 1$, together the bounds: $|w_{t+1}(\sigma)| \in [\underline{w}, \bar{w}]$, with $\underline{w}$ and $\bar{w}$ uniform in t . The feasibility of the upper bound $\bar{w}$ here follows from the geometric convergence in the product ratio to zero, and the possibility of satisfying a lower bound is trivial. Given these values for the sequence $\{w_t(\sigma)\}_{t \geq \tau}$, set ε_t to satisfy:

$$\varepsilon_t = \gamma |\lambda_t^k(\sigma) w_t(\sigma)|$$

The bounds on $\lambda_t^k(\sigma)$ and $w_t(\sigma)$ imply ε_t is bounded above zero. Using this in (62), we have:

$$\begin{aligned} \beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right] w_{t+1}(\sigma) &\geq (1 + \gamma) |\lambda_t^k(\sigma) w_t(\sigma)| \\ &= |\lambda_t^k(\sigma) w_t(\sigma)| + \varepsilon_t \\ &\geq \lambda_t^k(\sigma) w_t(\sigma) + \varepsilon_t \end{aligned} \quad (64) \quad (65)$$

If true for a positive-measure subset of $\sigma \in \Sigma$, this would imply a strict improvement for all $t \geq \tau$ – contradicting that the policy is time-consistently undominated. A symmetric argument can be applied when the product ratio is exploding (see proof of Proposition 7 above) completing the proof. $\square$

An equivalent to Proposition 8 goes through for this case with only cosmetic adjustments to the proofs: a time-consistently undominated policy requires that the condition:

$$\frac{\lambda_{t+1}^k(\sigma)}{\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta}} = \alpha_t$$

is almost never violated at the limit as $t \rightarrow \infty$, where $\alpha_t \in (0, 1]$ for all t .

The equivalent of Proposition 9 is:

Proposition 15. *Consider a policy $(\bar{\mathbf{x}}'_s, \bar{\mathbf{a}}'_s)$ that solves Problem 1 for the promise sequence that it induces, $\bar{\omega}'_s$. The continuation of this policy $(\bar{\mathbf{x}}'_t, \bar{\mathbf{a}}'_t)$ will belong to $D(x'_{t-1})$ for all $t \geq s$ provided the following are true:*

1. *The value function $V(\bar{\omega}_s; x_{s-1})$ is concave in $\bar{\omega}_s$.*
2. (a) *There exist positive scalars $\underline{K}$ and $\bar{K}$ such that for all $\tau \geq s$, $r > \tau$ and $\sigma \in \Sigma$:*

$$\underline{K} \leq \prod_{t=\tau}^{r-1} \left| \frac{\lambda_t^k(\sigma)}{\beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right]} \right| \leq \bar{K} \quad (66)$$

- (b) *There is a sequence of scalars $\{\alpha_t\}_{t=s}^{\infty}$, with $\alpha_t \in [\underline{\alpha}, \bar{\alpha}]$ for all t and $0 < \underline{\alpha} \leq \bar{\alpha} < 1$, such that the multipliers converge across $\sigma \in \Sigma$ as follows:*

$$\lim_{t \rightarrow \infty} \left[\frac{\lambda_{t+1}^k(\sigma)}{\alpha_t \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right]} \right] = 1 \quad (67)$$

where the rate of convergence is at least linear.

The proof works identically to the proof of Proposition 9, substituting the object:

$$\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta}$$

for $\lambda_t^m(\sigma)$. Concavity of the value function for these problems will usually follow from the linearity of the forward-looking constraints in the within-period utility function, together with increasing marginal cost of providing utility.

Condition (66) implies that a policy satisfying the sufficiency conditions can be interpreted as a time-consistently optimal choice in a restricted-dimensional problem that allows period-by-period choice across promises. In particular, the policy must satisfy the

condition:

$$\begin{aligned}
0 = \sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma \in \Sigma} & \left\{ \beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) \right. \right. \\
& \left. \left. - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right] \delta_{t+1}(\sigma) \right. \\
& \left. - \lambda_t^k(\sigma) \delta_t(\sigma) \right\} d\Pi(\sigma)
\end{aligned} \tag{68}$$

for all $s \geq 0$, where $\{\delta_t(\sigma)\}_{t \geq s}$ is a bounded sequence of scalars for all $\sigma \in \Sigma$, with $\lim_{\tau \rightarrow \infty} \inf_{t \geq \tau} |\delta_t(\sigma)| > 0$ for Π -almost all σ . This corresponds to a within-period multiplier restriction:

$$\beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right] = \lambda_t^k(\sigma) \frac{\delta_t(\sigma)}{\delta_{t+1}(\sigma)} \tag{69}$$

A symmetric policy is defined as one that allows policymakers in all periods the same freedom to vary promises at the margin. This implies the simpler restriction:

$$\beta \left[\int_{\Theta} \lambda_t^m(\sigma_-, \theta, \tilde{\theta}) d\tilde{\theta} + \lambda_t^k(\sigma_-) - \int_{\Theta} \lambda_t^m(\sigma_-, \tilde{\theta}, \theta) \frac{\pi(\tilde{\theta})}{\pi(\theta)} d\tilde{\theta} \right] = \lambda_t^k(\sigma) \tag{70}$$

This is the condition used in the Atkeson-Lucas example in the main text (Section 9.3).

D Applications: further details of calculations

D.1 Capital taxation

D.1.1 Inner problem

The implementability condition (28) can be decomposed using promise values into promise-making and promise-keeping constraints, respectively:

$$u_{c^k,t}(c_t^k + k_t) \leq u_{c^k,t}c_t^k + \beta\omega_{t+1} \quad (71)$$

$$\omega_t \leq u_{c^k,t}c_t^k + \beta\omega_{t+1} \quad (72)$$

The inner problem in period s is to maximise W_s subject to (26), (71) and (72) holding for all $t \geq s$, given k_{s-1} and $\bar{\omega}_s$. First-order conditions for this problem with respect to c_t^w , c_t^k and k_t in turn are:

$$u_{c^w,t} - \eta_t = 0 \quad (73)$$

$$\mu u_{c^k,t} - \eta_t - \lambda_t^m u_{c^k c^k,t} k_t + \lambda_t^k [u_{c^k,t} + u_{c^k c^k,t} c_t] = 0 \quad (74)$$

$$-\eta_t + \beta\eta_{t+1} [1 + f_{k,t+1} - \delta] - \lambda_t^m u_{c^k,t} = 0 \quad (75)$$

where η_t is the multiplier on the resource constraint (26).

D.1.2 Ramsey policy

Ramsey policy for period 0 is characterised by the first-order multiplier conditions:

$$\lambda_0^k = 0 \quad (76)$$

$$\lambda_t^k = \lambda_{t-1}^k + \lambda_{t-1}^m \quad (77)$$

for $t > 0$. Using these in (74) and (75) delivers a system of dynamic equations studied by Straub and Werning (2015). As these authors show, for $\sigma > 1$ the result is for the capital stock to converge to a ‘corner’ solution. When μ is sufficiently small, this involves zero long-run consumption for workers, with just sufficient capital to ensure government expenditure is sustained.

D.1.3 Time-consistently undominated policy

Given ratio comparability, a symmetric time-consistently undominated policy implies the condition:

$$\lambda_t^k \omega_t = \beta (\lambda_t^k + \lambda_t^m) \omega_{t+1} \quad (78)$$

for all t . Combining this with conditions (71) and (72), together with complementary slackness, reduces it to:

$$\lambda_t^k c_t^k = \lambda_t^m k_t \quad (79)$$

This allows λ^k and λ^m to be eliminated from (74) and (75), which collapse to the single condition:

$$k_t \{ \beta \eta_{t+1} [1 + f_{k,t+1} - \delta] - \eta_t \} = c_t^k \{ \eta_t - \mu u_{c^k,t} \} \quad (80)$$

This is equation (29) in the main text.

D.2 Limited commitment

D.2.1 Inner problem

The forward-looking constraints can be decomposed into the following two promise-making restrictions:

$$u(c_t(\sigma)) + \beta \omega_{t+1}(\sigma) \geq V(\sigma) \quad (81)$$

$$u(c_t(\infty)) + \beta \omega_{t+1}(\infty) \geq \frac{u(y^l)}{1 - \beta} \quad (82)$$

where (82) is for $\sigma > 0$, and the following two promise-keeping restrictions:

$$\mathbb{E}_{t-1} [u(c_t(\sigma')) + \beta \omega_{t+1}(\sigma') | \sigma] \geq \omega(\sigma) \quad (83)$$

$$u(c_t(\infty)) + \beta \omega_{t+1}(\infty) \geq \omega_t(\infty) \quad (84)$$

The time-consistent inner problem in period s maximises W_s subject to (32), and (81) to (84), given the sequence of state-contingent promises $\bar{\omega}_s$. Provided the utility function is concave, first-order conditions are necessary and sufficient for this. Normalising the multipliers for population sizes, we require, for all t and all σ :⁴⁹

$$u'(c_t(\sigma)) (1 + \lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)) - \eta_t = 0 \quad (85)$$

where η_t is again the resource multiplier. This is a standard optimality condition for a cross-sectional allocation problem, with $(1 + \lambda_t^m(\sigma) + \lambda_t^k(\sigma_-))$ the effective Pareto weight on an agent of type σ . The only departure from a first-best allocation is that Pareto weights may be changing over time for a given individual.

⁴⁹When $\sigma = 0$ the predecessor σ_- may take on many values, and the condition could be rewritten to allow for this by aggregating across corresponding values of $\lambda_t^k(\sigma_-)$. However in practice this is precisely the case in which the promise-making constraint binds, for both the Ramsey and TCUP solutions. This means the combined Pareto weight $(1 + \lambda_t^m(\sigma) + \lambda_t^k(\sigma_-))$ always takes the same value for $\sigma = 0$, irrespective of σ_- . See below.

D.2.2 Ramsey policy

Ramsey policy is characterised by the multiplier recursions:

$$\lambda_0^k(\sigma) = 0 \quad (86)$$

$$\lambda_t^k(\sigma) = \lambda_{t-1}^m(\sigma) + \lambda_{t-1}^k(\sigma_-) \quad (87)$$

for all $\sigma \in \Sigma$.

D.2.3 Time-consistently undominated policy

Under utilitarianism utility can be assumed to be difference-comparable, so that symmetric time-consistently undominated policy replaces the Ramsey multiplier condition with:

$$\lambda_t^k(\sigma) = \beta [\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)] \quad (88)$$

for all $\sigma \in \Sigma$.

Again, the solution has the property that promise-making constraints only bind for agents with $\sigma = 0$. This, together with the discounting and timing structure of (88), allows the Pareto weight to be rewritten as:

$$1 + \lambda_t^m(\sigma) + \lambda_t^k(\sigma_-) = 1 + \beta^\sigma \lambda_t^m(0) \quad (89)$$

D.3 Asymmetric information

D.3.1 Inner problem

Constraint (38) can be decomposed into ‘promise making’ and ‘promise keeping’ components. The promise making constraint is:

$$\theta^l u(c_t(\sigma_-, \theta^l)) + \beta \omega_{t+1}(\sigma_-, \theta^l) \geq \theta^l u(c_t(\sigma_-, \theta^h)) + \beta \omega_{t+1}(\sigma_-, \theta^h) \quad (90)$$

for all $\sigma_- \in \Sigma$.⁵⁰ The promise keeping constraint is:

$$\mathbb{E}_{t-1} [\theta u(c_t(\sigma_-, \theta)) + \beta \omega_{t+1}(\sigma_-, \theta) | \sigma_-] \geq \omega_t(\sigma_-) \quad (91)$$

where expectations are taken across period- t θ draws.

The inner problem is to maximise (36) subject to (37), (90) and (91). The multiplier on (90) is denoted $\lambda_t^m(\sigma_-, \theta^l)$, consistent with the shock history of agents for whom it binds. First-order conditions for this problem with respect to $c_t(\sigma_-, \theta^h)$ and $c_t(\sigma_-, \theta^l)$

⁵⁰The notation (σ_-, θ) denotes the history σ_- followed by θ . Replacing (38) with this constraint exploits the one-shot deviation principle.

in turn are:

$$u_{c,t}(c_t(\sigma_-, \theta^h)) \theta^h \left\{ 1 + \lambda_t^k(\sigma_-) - \lambda_t^m(\sigma_-, \theta^l) \frac{\theta^l(1-p)}{\theta^h p} \right\} - \eta_t = 0 \quad (92)$$

$$u_{c,t}(c_t(\sigma_-, \theta^l)) \theta^l \{ 1 + \lambda_t^m(\sigma_-, \theta^l) + \lambda_t^k(\sigma_-) \} - \eta_t = 0 \quad (93)$$

This can again be interpreted as the solution to a cross-sectional allocation problem in which Pareto weights for the different types are given by the objects in curly brackets. Optimal choice of assets through time implies a standard Euler condition:

$$\eta_t = \beta R \eta_{t+1} \quad (94)$$

Useful insight into the character of the solution is obtained by combining (92) and (93) to yield:⁵¹

$$\eta_t \mathbb{E}_{t-1} \left[\frac{1}{u_{c,t}(c_t(\sigma_-, \theta))} \right] = 1 + \lambda_t^k(\sigma_-) \quad (95)$$

D.3.2 Ramsey policy

A Ramsey-optimal choice implies the following conditions for the promise multipliers:

$$\lambda_t^k(\sigma_-, \theta^h) = -\lambda_{t-1}^m(\sigma_-, \theta^l) \frac{1-p}{p} + \lambda_{t-1}^k(\sigma_-) \quad (96)$$

$$\lambda_t^k(\sigma_-, \theta^l) = \lambda_{t-1}^m(\sigma_-, \theta^l) + \lambda_{t-1}^k(\sigma_-) \quad (97)$$

for all $\sigma_- \in \Sigma$, together with the normalisation:

$$\lambda_{-1}^m(\sigma) = \lambda_{-1}^k(\sigma) = 0 \quad (98)$$

for all $\sigma \in \Sigma$. Combining (96) and (97) gives:

$$\mathbb{E}_t [\lambda_{t+1}^k(\sigma_-, \theta)] = \lambda_t^k(\sigma_-) \quad (99)$$

That is, the promise-keeping multiplier follows a martingale process. Using this in (95), together with (94), gives the inverse Euler condition (39) in the main text.

D.3.3 Time-consistently undominated policy

In this problem promises correspond to utility values, and since the policymaker is utilitarian these must be difference comparable across individuals with different σ draws.

⁵¹This uses the normalisation $\bar{\theta} = 1$, and the independence of $\lambda_t^k(\sigma_-)$ with respect to the period- t shock.

The extension of the symmetric multiplier condition (23) to this case implies:⁵²

$$\lambda_t^k(\sigma_-, \theta^h) - \beta \left[-\lambda_t^m(\sigma_-, \theta^l) \frac{1-p}{p} + \lambda_t^k(\sigma_-) \right] = 0 \quad (100)$$

$$\lambda_t^k(\sigma_-, \theta^l) - \beta [\lambda_t^m(\sigma_-, \theta^l) + \lambda_t^k(\sigma_-)] = 0 \quad (101)$$

Replacing the Ramsey conditions (96) and (97) with these delivers the symmetric time-consistently undominated solution. A useful contrast is obtained by combining (100) and (101) to give:

$$\mathbb{E}_{t-1} [\lambda_t^k(\sigma_-, \theta)] = \beta \lambda_t^k(\sigma_-) \quad (102)$$

Thus promise-keeping multipliers decay at rate β in expectation along a sample path for past type draws, but again this is true *within a given period* t . The cross-sectional equivalent of the inverse Euler equation, condition (40), can then be obtained, using (95).

⁵²See Appendix C