Time-Consistently Undominated Policies*

Charles Brendon[†] Martin Ellison[‡]

September 15, 2017

Provisional, unedited draft -- comments welcome

Abstract

This paper reconsiders the problem of normative policy design in classic Kydland and Prescott problems, allowing for a commitment device. By definition, no policy in these environments can be optimal from the perspective of each successive period. The standard normative response is to commit to policies that are optimal from the perspective of a single initial period. An alternative, explored here, is to find a weaker selection criterion than optimality that *can* be satisfied at every point in time. To this end, we define an incomplete 'dominance' ordering across alternative options, ruling out choices that violate either a dynamic version of the Pareto principle or a simple independence criterion. Given this, we derive necessary and sufficient conditions for policies to be time-consistently undominated, and show that these can be satisfied via a straightforward change to the first-order conditions that apply under Ramsey-optimal policy. The resulting policies exhibit simpler initial dynamics than Ramsey policy, and significant long-run differences. Applications include a New Keynesian inflation bias problem, a Chamley-Judd capital tax problem, and a social insurance problem with limited commitment.

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, 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. All errors are ours.

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

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

1 Introduction

Time inconsistency is an endemic problem in the macroeconomic policy literature. Whether monetary, fiscal, 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 any particular time period depends on when this choice is being assessed – is it best *ex-ante*, or contemporaneously? The implied inconsistency in optimal choice was famously formalised by Kydland and Prescott (1977), and its consequences have been widely studied by macroeconomists from both a normative and a positive perspective 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 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 would be the equilibrium outcome of period-by-period, discretionary policy choice?* Discretionary choice, without any formal commitment device, can severely limit the efficiency of policy design, because credible promises cannot be made. The welfare implications are generally sensitive to the equilibrium concept used, but the simplest Markovian equilibria usually imply a lower welfare level in every period than could be attained through a feasible commitment. Normatively, they are not defensible selections. This is the well-known 'rules beat discretion' conclusion of Kydland and Prescott (1977).[1]

This positive analysis of discretionary choice is what justifies a commitment of some kind. Our focus is on how commitment should be designed. Thus the difference between our approach and Ramsey analysis comes in the method of choosing, not the basic set of constraints that choice faces. We are asking what normative criteria choice can satisfy, if it is required to satisfy these same

[1]History-contingent reputational equilibria may raise welfare by comparison, though the extent is unclear as these equilibria are generally subject to multiplicity. A fuller discussion is provided in Section 11 below.

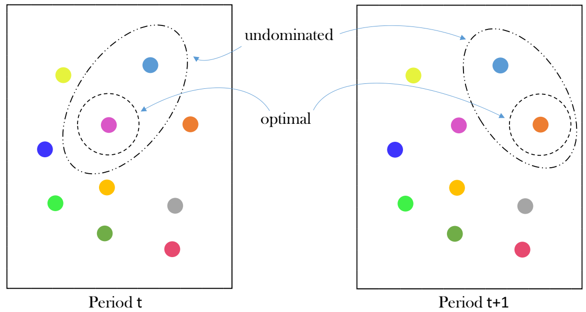

Figure 1: Time-consistently undominated policies: a heuristic illustration

criteria in each period successively.

Conceptually this is a difficult problem to pose, because choice and optimisation are so intrinsically connected. If a chosen policy is not to be *optimal* in the set of feasible options, how else can it be justified? The analytical device we use to confront this is an idea of dominated selections. Even when a general choice problem is subject to time inconsistency, some policy comparisons are less contentious than others. As an example, it might be possible to isolate a subset of the available options, and discover that for choice in this subset alone, no time inconsistency problem exists. If this is true, consistency suggests that a sub-optimal choice in the restricted subset would not be a desirable selection for the problem as a whole. Alternative options exist that are time-consistently superior.

Through variants on this form of reasoning, we show that a large number of feasible continuations may be ruled out as time-consistently dominated in any given period. Those that are not, are undominated. The set of undominated allocations will be larger than the more exclusive set of optimal choices, which will be contained within it. Time-consistent membership of the larger 'undominated' set may be possible where time-consistent membership of the optimal set is not. This is the basic normative argument that we pursue.

Figure 1 provides a heuristic illustration. Each coloured dot represents a possible dynamic allocation from period t onwards. The optimal allocation for period t, in pink, is not optimal in continuation in $t+1$. Thus there is time inconsistency in optimal choice. But when the larger 'undominated' set is considered, time-consistent membership is possible. The pale blue allocation is undominated in both periods, even though it is optimal in neither. The contribution of our paper is to present appealing dominance criteria that can be used to make a time-consistent normative selection in exactly this way.

1.1 Motivation

The motivation for developing this approach ultimately derives from unease expressed in the literature about the Ramsey benchmark, and the scope that a time-consistent normative method could have to overcome this. There are at least four distinct arguments that could be advanced, and we present them in turn.

1.1.1 The arbitrariness of date-contingent choice

As noted, Ramsey policies have been extensively studied and characterised across a range of important economic environments. A common feature, intrinsically connected to time inconsistency, is their date-contingent character. That is, instrument choice depends on the amount of time that has elapsed since the initial optimisation period. Optimal capital taxes, for instance, will often start at extremely high levels, but tend towards zero in the long run. An optimal monetary policy may involve a significant bout of initial inflation, but long-run price stability. Optimal social insurance might mean substantial initial redistribution, but ever-greater inequality as time progresses. Crucially, such trends are distinct from any evolution in the underlying structure of the economy over time. They arise simply because long-run policy is influenced by a need to make good on past promises, whereas initial choices are free from these constraints.[2]

A number of authors have argued that time variation in economic policy, independent of trends in the underlying economy, is either undesirable, implausible, or both. This view has been particularly prominent in the New Keynesian monetary policy literature, where criticism has focused particularly on its apparent arbitrariness. Svensson (1999), for instance, asked: "What is special about date zero?", whilst Woodford (1999) likewise challenged the idea of "being bound at all times by a commitment that appeared desirable at one single point in time". Instead, Woodford (2003) argued that what is needed is a "systematic *decision procedure* in the light of which ... current actions are always to be justified".[3] Defining such a procedure is a natural route to a choice that is date-invariant. If a policy can be justified successively according to the same set of principles, date zero ceases to be special.

Woodford (1999) goes on to recommend one possible candidate for a time-consistent decision procedure, which he calls 'optimality from a timeless perspective'. This essentially involves implementing in all time periods the long-run outcome that arises under Ramsey choice – discarding the short-term transition. This has come to be widely applied in the New Keyenesian policy literature, but its normative foundations are certainly questionable. As Woodford (2003) notes, it is a policy that would have emerged under an optimal plan devised in the distant past.[4] But is this an appropriate normative criterion to be applying? An important motivation for our paper is to provide a choice-theoretic framework for answering this.

[2]Section 2 illustrates these dynamics in a well-known inflation bias example.
[3]Woodford (2003), § 7.1, p. 474.
[4]See, for instance, Woodford (2003), § 8.1, p. 539.

1.1.2 Preserving desirability in the long run

A distinct set of objections to Ramsey policy has emerged in the social insurance literature. Here the issue is not so much the arbitrariness of treating some periods differently from others, but the specific fact that long-run Ramsey outcomes can be extremely undesirable when viewed in isolation.[5] The most famous example of this comes in the so-called 'immiseration' results from dynamic social insurance models. In dynamic Mirrlees and dynamic moral hazard settings, an optimal policy from the perspective of date zero will generally involve structuring incentives over time. Individuals who are lucky in early time periods will receive a persistent boost to their effective wealth, and individuals who are unlucky will see a persistent reduction. Under seemingly conventional assumptions on preferences and technology, it is possible to show that this results, under the Ramsey plan, in the consumption of almost all agents being driven to zero as time progresses – whilst a lucky subset of measure zero prosper. This arises even when the policy objective is utilitarian. The *ex-ante* incentive gains from extreme long-run inequality happen to outweigh the *ex-post* losses.

The deeper issue here is that an optimal choice for date zero alone need not have any particularly desirable properties when viewed in continuation at a later point in time. A time-consistent normative choice method ought to overcome this by design. Provided complete inequality is dominated by some alternative, it will never come to be implemented.

To date, the normative literature engaging with the immiseration result has focused on a very different strategy for overcoming it. In particular, Phelan (2006) and Farhi and Werning (2007) have shown that a change to the social discount factor, making the Ramsy planner more patient, can ensure a more benign long-run outcome. This is a perfectly defensible approach, in a normative tradition dating back to Ramsey (1928), but it has implications for policy design even in environments where time inconsistency is not an issue. A more patient social planner will generally wish to drive up the savings rate even in an efficient economy that satisfies the welfare theorems. A strength of our method will be to ensure that choice remains meaningfully desirable in the long run, without changing policy prescriptions where no time inconsistency exists.

1.1.3 Allowing flexibility

There are also more practical reasons why a time-consistently justifiable policy might be useful. In particular, it can offer more flexibility when the structure of the economy is uncertain. A Ramsey policy is a time-varying path of instrument choices from date zero onwards, chosen from the perspective of date zero. Suppose that at date 60, say, it is discovered that the economy's structure is different from what was first assumed. How should the Ramsey plan be updated? Should choices now be completely re-set, so as be optimal from date 60 onwards, given the new information? If the new information is relatively trivial,

[5]There is a clear tension with the timeless perspective's focus on long-run Ramsey outcomes here.

this would seem to violate the original commitment. Should instruments instead be set in a way that would have been optimal from the perspective of date zero, had the new information been known then? This feels artificial, and may not even be feasible, given that choices from periods zero to 59 have not been consistent with this plan. Should the past commitment simply be respected regardless? This is surely undesirable. The choice is not straightforward.

A time-consistently undominated selection can sidestep these issues. As a policy, it is equally defensible in date 60 as in date zero. If the structure of the economy has changed, the set of undominated policies in period 60 will likewise change, and a new selection can be made accordingly. There is no dependence on a past period's viewpoint to complicate matters.

1.1.4 A simple commitment

A final, more informal motivation relates to the benefits of simplicity in policy design, particularly when seeking to influence expectations. More or less by definition, a date-contingent policy cannot be summarised in a simple, universal rule of best conduct. An optimal capital tax plan may imply zero taxes in the long run, but only after a high-tax transition. Promising 'zero taxes, but not yet' is not the most direct route to fixing expectations. Likewise, commiting to a time-varying path for inflation lacks the simple focus of a single inflation target. Time inconsistency arises because changes to expectations can have real effects. Although we will work in rational expectations models where commitments are perfectly credible, in practice a promise to apply the same simple policy rule in every period is likely to have a far greater impact on expectations than publication of a date-contingent optimal plan. Again, finding a time-consistently applicable choice procedure is a necessary step towards finding time-invariant, desirable simple rules.

1.1.5 Adding to the toolkit

Though we believe that all of these arguments have some force, ultimately their validity is not our main concern. So long as reasonable doubts can exist about the appropriateness of Ramsey-optimal choice, it makes sense as a practical matter to investigate what time-consistent normative alternatives might be available. At the very least it is conceivable that certain governments and central bankers may wish to avoid proposing date-contingent policy.[6] Given this, a thorough analysis of alternative options can only enrich the policy toolkit.

1.2 Overview of approach

This subsection provides a more detailed sketch of our argument and main results.

[6]Arguably this is evidenced by the popularity of Woodford's 'timeless' choice method in the central banking literature.

1.2.1 Two dominance criteria

Identifying a set of undominated allocations period-by-period clearly requires a well-defined concept of dominance. Dominance is formalised in our paper as a binary relation on the space of feasible dynamic choices that are available in continuation from any given period t onwards. This binary relation will be incomplete, i.e. it will not rank every pair of feasible options, but where it exists it will coincide with the period-t policymaker's standard preference ordering across these allocations. That is, if allocation A strictly dominates allocation B in period t, then the policymaker in t strictly prefers A to B. Incompleteness, however, means that the converse need not be true. This incompleteness means that the set of undominated allocations from the perspective of t can be bigger than the set of optimal allocations.

The guiding principle that we use in constructing the dominance relation is to isolate time-consistent policy judgements, to the extent that these exist. This is done in two distinct ways. The first is based on finding **time-consistent constraint sets**. Suppose that a subset of the feasible allocations has the property that if choice were restricted to this subset alone, ignoring all other options, then standard optimisation would be time-consistent. Then inferior allocations in this subset will be judged to be dominated. This can be interpreted as a version of the independence of irrelevant alternatives condition, in the spirit of Nash (1950). If an available allocation would not be chosen when the constraint set is limited to allow time-consistent choice, then – under to this criterion – it should not be chosen in a larger problem, regardless of whether time inconsistency now arises.

In practical terms, this can be applied when comparing allocations that are identical in terms of the promise values that they deliver through time. Time inconsistency arises because of the incentive to make promises that are not subsequently kept. In all of the problems that we consider, these promises can be summarised as a requirement that some net-present value object – an agent's lifetime utility, say, or net expenditure – should attain a certain minimum value. This is labelled a 'promise value', following Abreu, Pearce and Stachetti (1990). Our requirement is equivalent to asserting that a chosen allocation should deliver its promise values in the most efficient way possible. Holding constant the promise values through time, there is no remaining time inconsistency.

The second criterion for imposing the dominance relation will be based on **time-consistent preference orderings**. Dominance is assumed to respect a variant of the Pareto principle, applied over time. Suppose that two allocations, A and B, are feasible in continuation at every point in time from t onwards.[7] We assert that A strictly dominates B in period t if there is a strict policy preference for the continuation of A over the continuation of B at every point in time from t on. Thus an allocation will never be included in the undominated set in t if there is a way to make all policymakers from t onwards strictly better

[7] Strictly, that the continuation of either allocation is always available, provided one or other has been pursued up to the current date. Thus pairwise policy preferences between A and B at each date relate to a non-hypothetical choice.

off. Among others, this will rule out the inefficient outcomes that can follow under discretion.

1.2.2 Time-consistently undominated choice is time-consistently optimal

Based on these criteria, together with a supplementary independence of irrelevant alternatives requirement, we proceed to characterise time-consistently undominated choice. To do this, we first highlight that the problem of choosing an undominated selection can be divided into a two-stage procedure. First, an undominated choice must solve an 'inner problem', delivering an optimal allocation for a given sequence of promise values over time. Second, the promise values themselves must be chosen in a way that implies an undominated choice for the inner problem.

Importantly, we show that this second step is equivalent to finding a sequence of promise values that is time-consistently *optimal* in a restricted choice set that places cross-restrictions on the promises that can be chosen at each point in time.[8] Heuristically, a time-consistently undominated policy finds promises that are optimal along a single restricted choice dimension from the perspective of every successive period. This contrasts with Ramsey policy, where promises are optimal along every choice dimension from the perspective of a single period – the first.

This link to an optimisation problem is particularly useful for practical purposes. It means that time-consistently undominated policies can be characterised by reference to simple first-order conditions. Satisfying these conditions is necessary and, with some tightening of assumptions, sufficient for finding a time-consistently undominated policy.

The first-order conditions that we derive are restrictions linking the shadow costs and benefits of keeping and making promises over time. They can be viewed as counterparts for time-consistently undominated policy to the well-known promise-multiplier restrictions that characterise Ramsey policy, as exploited by Marcet and Marimon (1998, 2016) to obtain a recursive saddle point representation of Ramsey choice. Analytically, this change to the multiplier restriction is the only difference between Ramsey policy and ours. Moreover, since the restrictions that we propose imply simpler dynamics, if anything the result is that time-consistently undominated policy – for all its set-theoretic foundations – is less demanding to compute than the Ramsey benchmark.

[8] That is, in the restricted problem it may be possible to amend a promise that was to be kept in period t, but only in conjunction with a simultaneous change in equivalent promises for $t+1$, $t+2$, and so on. A concrete example is a requirement that a chosen inflation target should be constant through time. Any increase in the current target is possible only alongside a rise in the rate that is promsied for the future. Section 2 explores a version of the classic inflation bias problem by way of motivation.

1.2.3 Long-run Ramsey policy is time-consistently dominated

Significantly, we prove that the necessary conditions for a time-consistently undominated policy are systematically violated by Ramsey policy. In particular, the long-run continuation of Ramsey policy is generally inferior to alternative feasible choices, from the perspective of all current and future time periods. This point is illustrated under general theoretical assumptions, and in specific computed examples. It matters because it implies that the long-run Ramsey outcome should not be viewed as a desirable selection in itself, independently of the transition to it. It is the best plan to choose for period zero, but not the best plan to end up inheriting in the long run. Yet the influential 'timeless perspective' approach of Woodford (1999, 2003) recommends implementing the long-run Ramsey outcome in every period, omitting any transition. Our results indicate that this approach is straightforwardly dominated by alternative feasible selection criteria. Another implication of the result is that a zero long-run capital tax rate is only ever likely to be justifiable as part of a dynamic plan that involves very high initial capital taxes. As a time-invariant policy recommendation, it is likely to be dominated.

Intuitively, these results arise because a Ramsey policy does not just make sure that choices are on the Pareto frontier. It also ensures that the gains from moving to the frontier accrue as far as possible in the very first time period. To some extent, this comes at the expense of welfare in subsequent periods, when allocations are forever skewed towards supporting prior expectations. Time-consistently undominated policy does not allow this 'hangover' to arise. Policy must remain on the Pareto frontier in continuation in every period.

1.2.4 Multiplicity and symmetry

A complication to the argument is that time-consistently undominated policy is not generally unique. The sufficiency conditions that we derive can generally be satisfied by a continuum of different choices. Fortunately, a simple symmetry refinement is enough to eliminate this indeterminacy. Unfortunately, the exact meaning of symmetry will be example-specific.

Intuitively, symmetry means that the set of equations that characterise time-consistently undominated policy will take the same mathematical form in every time period. In particular, the key promise multiplier restriction will not be arbitrarily time-varying. Since the purpose of our method is to arrive at a choice technique that will be consistent through time in its selections, avoiding arbitrary time variation is an uncontroversial principle to adopt. Yet even this can be done in more than one way. Making progress requires a more focused, example-specific understanding of what exactly the promises represent, and what invariance properties their choice should satisfy. As an example, if promises relate to utility, and utility in each period is only defined up to the addition of an arbitrary scalar, it would not make sense to define symmetry in a way that was particular to one possible normalisation only. This means, for instance, that we define symmetry slightly differently in models where promises relate to utility, as against models where they relate to wealth levels. Though the difference makes for a slightly more involved presentation, it ultimlately ensures more coherent policy statements.

2 Motivating examples

The motivation for our paper can be clarified by exploring a simple example. This section explains the problem of time-invariant normative choice in the context of a linear-quadratic New Keynesian inflation bias problem with no uncertainty. With just two variables and one linear constraint, the purpose is to keep the environment as simple as possible.

2.1 Setup

We first consider a well-known New Keynesian variant on the classic inflation bias problem, with no uncertainty. 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:

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

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 β and γ 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 $\bar{\mathbf{y}}_\mathbf{0} := \{y_t\}_{t=0}^{\infty}$, $\bar{\boldsymbol{\pi}}_\mathbf{s} := \{\pi_t\}_{t=s}^{\infty}$, and so on. Finite sequences are distinguished by an additional superscript to denote end period: $\bar{\mathbf{y}}_\mathbf{0}^\mathbf{s} := \{y_t\}_{t=0}^{s}$, and so on.

2.2 The feasible set

Any pair $(\bar{\mathbf{y}}_\mathbf{s}, \bar{\boldsymbol{\pi}}_\mathbf{s})$ 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:

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

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. There are many aspects to this problem, including the question of whether any particular selection could be supported in a non-cooperative dynamic equilibrium – without a formal commitment device. Clearly this is an important positive question, but our main focus is normative, and we leave it aside. Thus we will allow in principle for any element of Ξ to be selected, under the assumption that commitment to this choice will be possible.

In any given period $s \geq 0$ the policymaker has preferences over the set of continuation policies $(\bar{\mathbf{y}}_{\mathbf{s}}, \bar{\pi}_{\mathbf{s}}) \in \Xi$, described by the loss function L_s:

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

where $y^* > 0$ is an optimal level for the output gap and χ is a parameter.[9] These preferences have the conventional additively separable form with geometric discounting, so are not a source of time inconsistency themselves.

The Ramsey policy is defined as the selection $\left(\bar{\mathbf{y}}_{\mathbf{0}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{0}}^{\mathbf{R}}\right)$ such that L_0 is minimised on Ξ. In this simple linear-quadratic environment it will be unique:

$$\left(\bar{\mathbf{y}}_{\mathbf{0}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{0}}^{\mathbf{R}}\right) = \arg \min_{(\bar{\mathbf{y}}_{\mathbf{0}}, \bar{\pi}_{\mathbf{0}}) \in \Xi} L_0$$

This policy choice is an important and widely-studied benchmark, but it is well known that it is a time-inconsistent selection. The time inconsistency is illustrated by Figure 2. This charts $\left(\bar{\mathbf{y}}_{\mathbf{0}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{0}}^{\mathbf{R}}\right)$ for a set of conventional parametric assumptions.[10] In initial periods inflation is relatively high, allowing output to be brought close to the ideal value y^*. As time progresses the policymaker delivers lower inflation, which is desirable from the perspective of period 0 as a way to contain initial inflation expectations. In the long run both π_t and y_t converge to zero. 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. Since the inflation-output paths in Figure 2 are time-varying, this clearly implies a different choice for s onwards from $\left(\bar{\mathbf{y}}_{\mathbf{s}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{s}}^{\mathbf{R}}\right)$. That is:

$$\left(\bar{\mathbf{y}}_{\mathbf{s}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{s}}^{\mathbf{R}}\right) \notin \arg \min_{(\bar{\mathbf{y}}_{\mathbf{s}}, \bar{\pi}_{\mathbf{s}}) \in \Xi} L_s$$

Thus the criterion 'minimise L_s on Ξ' is not a selection rule that can be applied consistently for all $s \geq 0$. This is the basic observation first due to Kydland and Prescott (1977). 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 expectations. When the future arises, the justification for keeping these promises has passed.

2.4 Time-consistent choice criteria

Time inconsistency presents a fundamental problem for normative policy design. There is no option in Ξ that is optimal to select initially, and optimal to continue

[9]It is well known that a desire to target a positive output gap can be motivated by uncorrected monopoly distortions in the product market, which restrict aggregate production below its efficient level in the underlying non-linear New Keynesian model. See, for instance, Benigno and Woodford (2005).

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

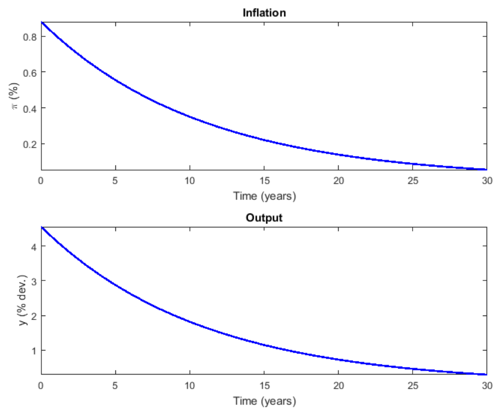

Figure 2: Inflation and output under Ramsey policy

with from the perspective of every period. A defensible selection from Ξ can at best either (a) satisfy optimality from the perspective of just one period, or (b) satisfy some weaker normative criterion than optimality, in every period.

Choosing the Ramsey plan $\left(\bar{\mathbf{y}}_{\mathbf{0}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{0}}^{\mathbf{R}}\right)$ means following the first approach, where period 0 is the date that is privileged. This option certainly has normative appeal. After all, if a commitiment device exists in period 0, why not commit to the policy that is best from this perspective? Yet a number of authors have expressed concern about the time-varying character of the resulting choice. Recall that the economic environment in this example is entirely stationary. There are no endogenous states or shocks that make, say, period 60 any different from period 0. And yet the Ramsey policy recommends a quite different inflation-output combination for period 60 compared to period 0. This arises entirely because period 0 is the period in which optimisation is permitted to take place – a designation that seems quite arbitrary. This in turn creates an impression of arbitrariness in the dynamics, succinctly summarised by Svensson's (1999) question: "What is special about date zero?"

Figure 3 reinforces this view. It charts the value of the policymaker's discounted loss function L_s over time, under the Ramsey policy. Despite the entirely stationary economic environment, the policymaker is strictly better off in period zero relative to all others. The chosen allocation is designed to minimise loss in the first period, even though the implications of this may be a quite different outcome subsequently. Clearly it is debatable whether this 'unique privilege' for the initial period is a desirable feature of policy design, particularly if some of the gains in period 0 come at the cost of higher-than-necessary loss in later periods. As we show below, this is indeed the case.

Moreover, there is a basic disparity between the time-varying form that Ramsey policy necessarily takes, and observed practice. Many central banks around the world *have* committed to an inflation target, with the aim of influencing

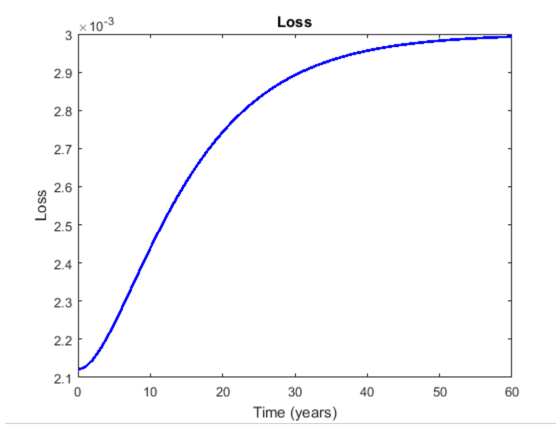

Figure 3: Loss under Ramsey policy

long-term inflation expectations. Yet none of these banks has specified their target to be time-varying, in a way that privileges the time period in which policy design took place. Woodford (2003) develops this point forcefully, arguing that the 'exceptional' character of the initial inflationary policy would be difficult to justify in practice, and may even undermine a central bank's ability to make credible commitments. Analysing the same problem, he writes that: *"[Ramsey policy] assumes that it is possible to commit to an arbitrary time path for inflation and have this be expected by the private sector; it is assumed to be possible to choose inflation 'just this time' while committing never to create inflation in the future. But there is reason to fear that the public should observe the central bank's method of reasoning, rather than its announced future actions, and conclude instead that in the present it should always wish to create inflation 'just this time'."*[11]

This motivates a consideration of the second possibility. If no policy is *optimal* from the perspective of every period, might there nonetheless be defensible choice criteria that can be repeatedly satisfied? Woodford (2003) poses exactly this problem when he argues for "a systematic decision procedure, in the light of which current actions are always to be justified".[12] Time inconsistency means that the procedure 'minimise $L_s$on Ξ' cannot satisfy this description. So what could? This problem is the main focus of our paper.

2.5 Optimal constant policy

The simplicity of the present example is helpful in making initial progress on this problem. It follows from the discussion above that a time-consistent choice procedure in a stationary, deterministic environment such as this must imply a

[11]Woodford (2003), Ch 7.1, pp 474.
[12]ibid.

constant outcome.[13] This is useful here because the class of constant inflation-output combinations is easy to investigate. Evidently this focus is only appropriate because of the particular simplicity of the current problem, but if we know the basic form that time-consistent choice must take in this example, it may be easier to find insights that can generalise.

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

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

Ξ^c is a particularly interesting subset of Ξ because all of its elements are time-invariantly ranked relative to one another. That is, the comparison between $(\bar{\mathbf{y}}'_{\mathbf{0}}, \bar{\pi}'_{\mathbf{0}}) \in \Xi^c$ and $(\bar{\mathbf{y}}''_{\mathbf{0}}, \bar{\pi}''_{\mathbf{0}}) \in \Xi^c$ under the loss criterion L_0 is always identical to the comparison between the respective continuations $(\bar{\mathbf{y}}'_{\mathbf{s}}, \bar{\pi}'_{\mathbf{s}}) \in \Xi^c$ and $(\bar{\mathbf{y}}''_{\mathbf{s}}, \bar{\pi}''_{\mathbf{s}}) \in \Xi^c$ under L_s. This is an immediate consequence of the allocations – and thus loss – being constant. An implication of this is that the normative problem of selecting from Ξ^c is 'timeless': if the set of options were restricted to Ξ^c, policy preferences across these options could be described independently of the time period in which the assessment takes place.

There is a unique choice that minimises L_s on Ξ^c for all s, which we label $(\mathbf{y}^{\mathbf{c}}_{\mathbf{0}}, \pi^{\mathbf{c}}_{\mathbf{0}})$ – the optimal constant policy.[14] A choice procedure that selected a different element of Ξ^c from $(\mathbf{y}^{\mathbf{c}}_{\mathbf{0}}, \pi^{\mathbf{c}}_{\mathbf{0}})$ would be extremely hard to justify normatively, since there is an alternative choice that is available every period and is superior under every period's loss criterion. Put differently, any normatively defensible time-consistent choice procedure ought to be recovering $(\mathbf{y}^{\mathbf{c}}_{\mathbf{0}}, \pi^{\mathbf{c}}_{\mathbf{0}})$ in this particular example.

2.6 Dominance and time consistency

This is a less trivial observation than it first seems, because it turns out that $(\mathbf{y}^{\mathbf{c}}_{\mathbf{0}}, \pi^{\mathbf{c}}_{\mathbf{0}})$ is not related in any obvious way to the most common policy benchmarks that exist in the literature. In particular, and perhaps surprisingly, it is not the limiting outcome that arises under Ramsey policy as time progresses. Figure 4 illustrates this, charting the values for inflation and output under the Ramsey plan and the optimal constant policy. The limiting Ramsey outcome involves strictly lower inflation and output levels than the optimal constant choice. By definition, when this is implemented as a constant outcome for all t, loss is strictly higher than under $(\mathbf{y}^{\mathbf{c}}_{\mathbf{0}}, \pi^{\mathbf{c}}_{\mathbf{0}})$. Figure 5 confirms this, charting the value of L_s over time under Ramsey policy and the optimal constant solution. Continuing with Ramsey policy becomes inferior to the optimal constant policy once sufficient time has elapsed. Hence the long-run Ramsey limit is strictly inferior as a constant choice.

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

[14] It is a simple exercise to show that this is given by $y^c_t = \frac{\chi(1-\beta)^2}{\gamma^2+\chi(1-\beta)^2} y^*$ and $\pi^c_t = \frac{\chi\gamma(1-\beta)}{\gamma^2+\chi(1-\beta)^2} y^*$ for all t.

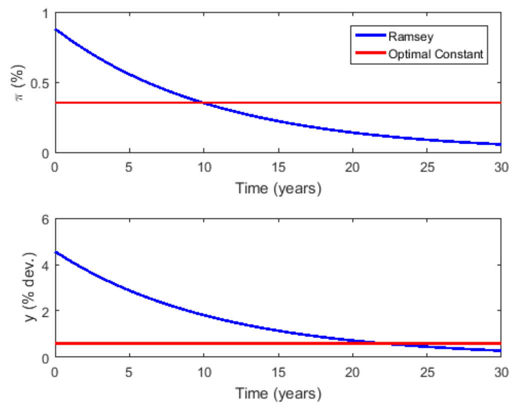

Figure 4: Ramsey and optimal constant policy

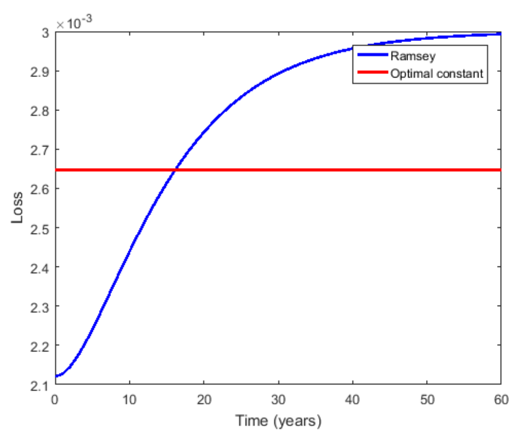

Figure 5: Ramsey and optimal constant policy: loss

In the context of the literature, these simple numerical observations matter because the long-run outcome of Ramsey policy is precisely what is advocated by Woodford (1999, 2003) as a desirable time-consistent choice. Labelling this "optimal from a timeless perspective", he argues that it is a policy "of the kind that one would always wish *to have been expected* to follow".[15] Yet, as this example highlights, the long-run limit of a policy path whose purpose is to maximise welfare in period 0 need not exhibit any particularly desirable properties *per se*. Inflation is kept low in the long run under Ramsey policy precisely so as to improve the inflation-output trade-off in earlier periods, given the forward-looking character of the Phillips curve (1). Woodford's timeless solution permits this improved trade-off, but never allows a prior period's decisionmaker to exploit it – a permanent hangover without the preceeding party. It is an allocation that is unambiguously inferior to the optimal constant choice.

Another notable constant policy benchmark that differs from $(\mathbf{y_0^c}, \pi_\mathbf{0}^\mathbf{c})$ in this model is the Markov equilibrium. This is the inflation-output combination that results from non-cooperative, period-by-period policy choice, under the assumption that choice is a function of payoff-relevant state variables only – which in this simple case means constant choice. The Markov equilibrium, denoted $\left(\bar{\mathbf{y}}_\mathbf{0}^\mathbf{d}, \bar{\pi}_\mathbf{0}^\mathbf{d}\right)$, is unique here, and exhibits the familiar inflation bias property. Its dynamics are charted in Figure 6, alongside the Ramsey and optimal constant policies.[16] It is an outcome that is commonly viewed as a coordination failure, and the loss comparison in Figure 7 confirms this. In all $s \geq 0$, a switch from $\left(\bar{\mathbf{y}}_\mathbf{s}^\mathbf{d}, \bar{\pi}_\mathbf{s}^\mathbf{d}\right)$ to $(\bar{\mathbf{y}}_\mathbf{s}^\mathbf{c}, \bar{\pi}_\mathbf{s}^\mathbf{c})$ would reduce the value of L_s. Though Markov choice may arguably be a *realistic* time-consistent choice procedure, it is hard to justify it as a desirable one.

A point that is worth emphasising here, however, is that the inefficiency of Markov equilibrium is qualitatively identical to the inefficiency of Woodford's timeless perspective policy. In both cases there is an alternative option that is available and strictly preferred in every current and future time period. Moreover, this inefficiency is also a feature of *the continuation* of Ramsey policy once sufficient time has elapsed. As Figure 5 makes clear, for a sufficiently large s, the value of the loss function L_t under the Ramsey continuation is greater for all $t \geq s$ by comparison with the optimal constant policy. So if the welfare properties of the Ramsey continuation were assessed in s, it could be described as inefficient for exactly the same qualitative reason as the Markov equilibrium.

We can build on the preceeding discussion by defining dominance across pairs of allocations as follows:

Definition. Policy $(\bar{\mathbf{y}}'_\mathbf{s}, \bar{\pi}'_\mathbf{s}) \in \Xi$ **dominates** the alternative $(\bar{\mathbf{y}}''_\mathbf{s}, \bar{\pi}''_\mathbf{s}) \in \Xi$ in period $s \geq 0$ if L_t is strictly lower under $(\bar{\mathbf{y}}'_\mathbf{t}, \bar{\pi}'_\mathbf{t})$ than $(\bar{\mathbf{y}}''_\mathbf{t}, \bar{\pi}''_\mathbf{t})$ for all $t \geq s$, including at any limit as $t \to \infty$.

This is a version of the familiar (strict) Pareto ordering, taken with respect to the preferences of policymakers at different points in time. As Figure

[15]Woodford (2003), Ch 7.1, pp 474, italics original.

[16]The analytical solution is $y_t = \frac{\chi(1-\beta)}{\gamma^2+\chi(1-\beta)} y^*$ and $\pi_t = \frac{\chi\gamma}{\gamma^2+\chi(1-\beta)} y^*$ for all t.

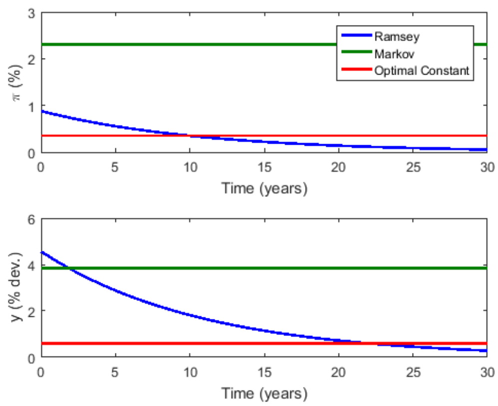

Figure 6: Ramsey, Markov and optimal constant policy

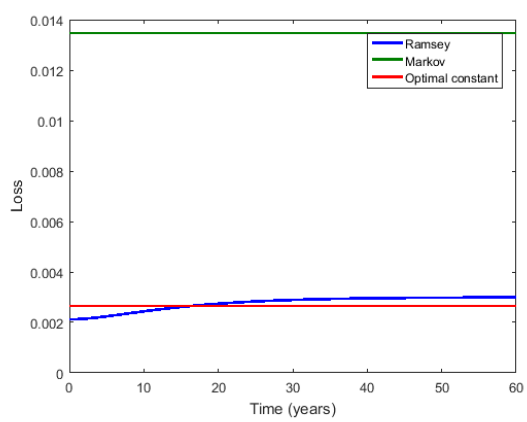

Figure 7: Ramsey, Markov and optimal constant policy: loss

7 makes clear, under this definition the Markov equilibrium $\left(\bar{\mathbf{y}}_{\mathbf{s}}^{\mathbf{d}}, \bar{\pi}_{\mathbf{s}}^{\mathbf{d}}\right)$ is dominated by both the optimal constant policy $(\bar{\mathbf{y}}_{\mathbf{s}}^{\mathbf{c}}, \bar{\pi}_{\mathbf{s}}^{\mathbf{c}})$ and the Ramsey continuation $\left(\bar{\mathbf{y}}_{\mathbf{s}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{s}}^{\mathbf{R}}\right)$ in all $s \geq 0$. In addition, as just discussed, for sufficiently large s $(\bar{\mathbf{y}}_{\mathbf{s}}^{\mathbf{c}}, \bar{\pi}_{\mathbf{s}}^{\mathbf{c}})$ comes to dominate $\left(\bar{\mathbf{y}}_{\mathbf{s}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{s}}^{\mathbf{R}}\right)$. But, like the usual Pareto criterion, dominance is not a complete relation. In period 0 neither $(\bar{\mathbf{y}}_{\mathbf{0}}^{\mathbf{c}}, \bar{\pi}_{\mathbf{0}}^{\mathbf{c}})$ nor $\left(\bar{\mathbf{y}}_{\mathbf{0}}^{\mathbf{R}}, \bar{\pi}_{\mathbf{0}}^{\mathbf{R}}\right)$ dominates the other.

A policy $(\bar{\mathbf{y}}_{\mathbf{s}}', \bar{\pi}_{\mathbf{s}}') \in \Xi$ 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 in addition. This is interesting because we know that in Kydland and Prescott problems, no single allocation inhabits the set of optimal choices indefinitely. By weakening the procedure 'select an optimal policy' to 'select an undominated policy', it may be possible to generate choices in a time-consistent, normatively appealing way. Accordingly, a policy $(\bar{\mathbf{y}}_{\mathbf{0}}', \bar{\pi}_{\mathbf{0}}') \in \Xi$ is described as **time-consistently undominated** if its continuation $(\bar{\mathbf{y}}_{\mathbf{s}}', \bar{\pi}_{\mathbf{s}}')$ is undominated for all $s \geq 0$. We then have the following result:[17]

Proposition 1. *The optimal constant policy* $(\bar{\mathbf{y}}_{\mathbf{0}}^{\mathbf{c}}, \bar{\pi}_{\mathbf{0}}^{\mathbf{c}})$ *is time-consistently undominated.*

The result is non-trivial, because we are saying that $(\bar{\mathbf{y}}_{\mathbf{s}}^{\mathbf{c}}, \bar{\pi}_{\mathbf{s}}^{\mathbf{c}})$ is undominated in the entire set Ξ, not just the restricted set Ξ^c in which it is optimal. It proves by example that time-consistently undominated policies can exist in environments where time-consistently optimal policies do not.[18] Note that the Ramsey policy, though undominated in period 0, is not time-consistently undominated. As Figure 5 illustrates, the optimal constant solution comes to dominate it after sufficient time has elapsed. The general lesson is that weakening the normative requirement from 'optimal' to 'undominated' allows for a choice criterion that can be asserted in all periods, at least in this example. Moreover, the implications for policy design appear to be non-trivial.

2.7 Multiplicity

An important complication, however, is that the optimal constant policy is not the *only* time-consistently undominated selection in this example. Figure 8 charts two others alongside it. The policy labeled 'limiting path', in blue, involves a strictly higher inflation rate initially, approaching the optimal constant choice at the limit as time passes. As Figure 9 illustrates, this path delivers a strictly lower value for loss than the optimal constant policy in every finite time

[17] Proofs of propositions are collected in the appendix.

[18] It is worth noting that the proof is dependent on the chosen definition of dominance. In particular, any attempt to weaken the definition of dominance so that strict preference was not required in all periods, including at the limit, would undermine the result. Intuitively, doing so would reduce the set of undominated policies in each period, to the point where no time-consistently undominated policy could exist.

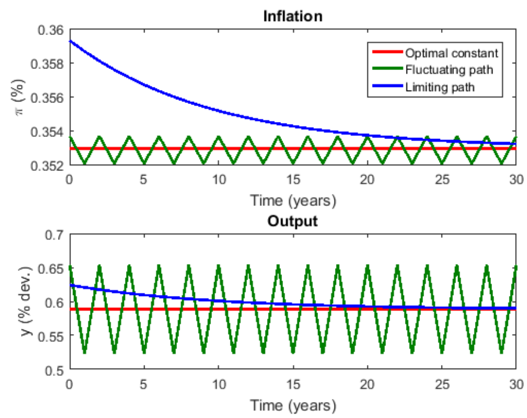

Figure 8: Three time-consistently undominated policies

period, but the two are identical by construction at the limit. Under the definition of dominance given above, limiting equivalence implies that neither option dominates the other – even though the blue path here is strictly preferred to $(\bar{\mathbf{y}}^{\mathbf{c}}_{\mathbf{s}}, \bar{\pi}^{\mathbf{c}}_{\mathbf{s}})$ for all finite s.[19] The example thus confirms that this limiting requirement is a non-trivial feature of the definition. Without it the optimal constant choice would be dominated, and thus so would all constant choices – effectively ruling out the possibility of achieving time-consistent normative selection by this route.

The other policy, labeled 'fluctuating path', in green, sees inflation and output follow a two-period cycle, permanently fluctuating about their optimal constant values. Every other period output and inflation are relatively high, and the value of loss falls below the level that is permanently attained under optimal constant policy. As Figure 9 illustrates, the policymaker in these alternate periods would prefer to adhere to the cyclical dynamics rather than switch to the optimal constant path. Proving that there is *no* alternative path that is preferred in all periods relies on results established later in the paper, and at present we simply assert it. What this example shows is that the multiplicity of time-consistently undominated solutions is not just a 'transition' issue. There are paths that do not converge to the optimal constant solution, whilst also being time-consistently undominated.

Two lessons can be drawn from these findings. The first, more negative, is that a normative theory of time-consistent choice cannot be based entirely on the idea of time-consistently undominated policies. At the very least, some additional strengthening will be needed to deliver a unique selection. Yet more positively, it is clear that one of the available options is simpler than all others. The optimal constant path is the only one that treats all time periods sym-

[19]More generally, a simple amendment to the proof of Proposition 1 will confirm that any allocation that limits to the optimal constant solution is time-consistently undominated.

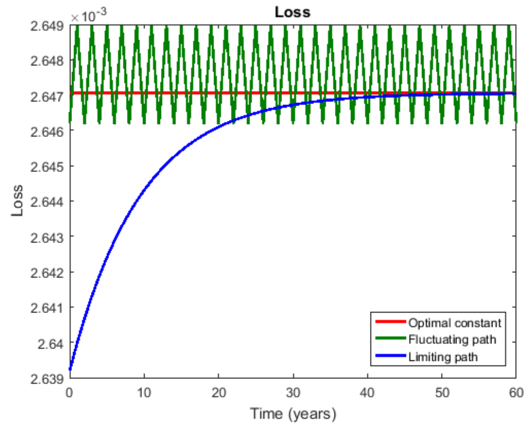

Figure 9: Three time-consistently undominated policies: loss

metrically. Generalising an idea of symmetry to richer environments – where the underlying economy is changing through time – is likely to pose further problems, but the broad possibility of a symmetry refinement seems to hold promise.

2.8 Summary

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

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 nether the outcome of a Markovian, discretionary equilibrium, nor the long-run outcome observed under Ramsey policy, where the latter corresponds to Woodford's (1999, 2003) 'timeless perspective' policy.

3. The optimal constant policy is time-consistently undominated.

4. Many other policies are also time-consistently undominated, but the optimal constant choice is the only one of these that is symmetric over time.

The analysis that follows will generalise all four of these insights. The idea of dominance that is used in this example extends easily to other settings, and we use it as one of a pair of dominance criteria that defines a broad set of undominated policies. Policies that are time-consistently undominated by these two criteria coincide neither with Markovian nor Ramsey/timeless choice. In general there are many time-consistently undominated policies, but an appropriate symmetry refinement can restrict focus to a singleton. Finally, we show that there is a very general link between time-consistently undomninated policies and time-consistently optimal choices in restricted sets of options – a generalisation of point 1. This link is particularly useful for practical purposes, because it allows the comparatively abstract idea of a time-consistently undominated policy to be linked to the solution to a relatively simple optimisation problem.

3 General setup

We develop these ideas 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 and superscripts end period where relevant. Thus $\bar{\mathbf{x}}_{\mathbf{s}} := \{x_t\}_{t=s}^{\infty}$, $\bar{\mathbf{x}}_{\mathbf{s}}^{\mathbf{r}} := \{x_t\}_{t=s}^{r}$, and so on.

3.1 Preliminaries

Time is discrete, and runs from period 0 to infinity. We neglect aggregate risk for simplicity, but the framework allows settings with idiosyncratic risk across large populations of agents. Adding aggregate risk is not difficult technically, but requires a stance to be taken on the social welfare weighting to give to alternative stochastic histories. This is a non-trivial consideration, and we opt to defer it.

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 the identifier σ varies flexibly across examples, but in general it is used to distinguish the complete 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 distinct forward-looking constraints, σ can be used as a simple index on these different constraints.

Importantly, we assume that Σ is a time-invariant set. In stochastic environments this means that even in period 0 there is assumed to be information about the prior realisation of shocks, and it is possible (though not necessary) for individual outcomes to differ according to these histories. There is no *a priori* requirement that this information should be used, and so allowing for its existence is not a restriction on choice.

σ is assumed to follow a Markov process over time, with the conditional probability measure $\Pi(S|\sigma)$ giving the probability of the subset $S \subseteq \Sigma$ in period $t+1$, given that σ is drawn in t.[20] Where the meaning is obvious, expectations with

[20] If the only role of σ is as a fixed index, this distribution becomes degenerate, with

respect to this measure will be represented by the standard operator $\mathbb{E}_t$. 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 knowledge of an individual's current σ draw implies knowledge their draw in all past time periods – so that σ is a rich identifier of shock history. Formally, this is done by assuming that σ is 'fully revealing' of past type:[21]

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

That is, σ' belongs to a set S of draws from Σ that can only be reached in the event that σ is the predecessor history. The assumption that we impose is that this property is satisfied by all $\sigma \in \Sigma$:

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

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

The problem in period s is to select a sequence of the form $(\bar{\mathbf{x}}_\mathbf{s}, \bar{\mathbf{a}}_\mathbf{s}) \in \mathcal{X} \times \mathcal{A}$, 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 each time period according to a given social preference ranking. Viewed in period s, this is described by the time-separable and σ-separable function W_s:

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

where $r : A_\sigma \times \Sigma \to \mathbb{R}$ is a within-period, σ-contingent preference function for period $s \geq 0$, and higher values of W_s correspond to more preferred outcomes. The direct dependence of r on σ will be redundant in many models, but in some

$\Pi(S|\sigma) = 0$ whenever $\sigma \notin S$, and $\Pi(S|\sigma) = 1$ whenever $\sigma \in S$.

[21] σ' is embedded in a larger set S here to allow for the possibility that σ' alone is a measure-0 event under $\Pi(\cdot|\sigma)$.

it is useful. For instance, in settings with idiosyncratic risk it allows the social objective to be a (weighted) utilitarian criterion across types of individuals.

These preferences are dynamically recursive, and so are not themselves a source of time inconsistency. Time inconsistency arises in Kydland and Prescott problems because of the nature of constraints, not preferences. 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.[22]

It is useful to define some 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 $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \succeq (\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ if and only if:

$$\sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma \in \Sigma} r\left(a'_t(\sigma), \sigma\right) d\Pi(\sigma) \geq \sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma \in \Sigma} r\left(a''_t(\sigma), \sigma\right) d\Pi(\sigma)$$

3.3 Constraints

Feasibility places two sets of constraints on the dynamics of x_t and a_t over time. First, there is an i-dimensional vector of 'structural' restrictions linking the inherited and future state vectors, and the controls:

$$p\left(x_{t-1}, x_t, a_t\right) \geq 0 \tag{4}$$

where $p : X \times X \times A \to \mathbb{R}^i$. 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 $K_t \leq (1-\delta) K_{t-1} + I_t$. When considering allocations from s onwards, condition (4) must be satisfied for all $t \geq s$.

Second, there is a set of infinite-horizon 'forward-looking' constraints, one for each $\sigma \in \Sigma$:

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

where $\sigma^{+\tau} \in \Sigma$ denotes a τ-period successor history to σ, $h : A_\sigma \times \Sigma \to \mathbb{R}$ and $h^0 : A_\sigma \times \Sigma \to \mathbb{R}$ for all $t \geq 0$. When considering allocations from s onwards, condition (5) must be satisfied for all $\sigma \in \Sigma$ at all $t \geq s$. The forward-looking character of this restriction generally implies time inconsistency, making it central to our analysis.

[22] For instance, in a model that features consumption habits it is possible that r might take the form $r\left(c_t - \lambda c_{t-1}\right)$ 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, as set out below.

3.3.1 Discussion of constraint (5)

The exact form of (5) reflects some important modelling choices that we now discuss.

Though it has considerable flexibility, the basic structure of the constraint most directly matches representations of incentive constraints in infinite-horizon environments, where h relfects a time-separable within-period payoff function, and h^0 provides a 'reservation' level for utility, wealth or equivalent. A well-known example is a utility-based participation constraint of the form $\mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^{\tau} u\left(c_{t+\tau}\right) \geq \bar{V}$. The parameter $\beta \in (0,1)$ in (5) can be interpreted as a standard discount factor. In the example of a participation constraint, the dependence of $a_{t+\tau}$ on σ' would capture the possibility that consumption in $t+\tau$ may depend on the complete history of earnings draws up to $t+\tau$. Notice that σ' is also allowed to affect the payoff function directly, as in the private information environment of Atkeson and Lucas (1992), for instance.

The infinite upper limit in the sum on the left-hand side of (5) is slightly restrictive, as it rules out some 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. Additionally, note that in some 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:

$$\pi_t = \gamma \mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^{\tau} y_{t+\tau} \tag{6}$$

When the equality is read as a two-sided inequality, this maps directly into (5), where dependence on σ has become trivial. The h^0 function is just π_t, and γy_t corresponds to h.

Finally, as with the objective function r, for simplicity we have assumed that state variables do not enter into h or h^0. This helps separate conventional feasibility restrictions from forward-looking constraints in the subsequent analysis. There is no obvious loss in generality, since again auxiliary variables and constraints can be used to capture any required state dependence.

3.4 Comparing h functions

Our main aim is to formulate a normatively desirable solution concept for Kydland and Prescott problems that is symmetrically applicable through time. Achieving this requires a characterisation of two ambiguous terms – desirability and symmetry. This subsection provides important foundations for an idea of symmetry.

Though the absence of symmetry can sometimes be easy to identify – particularly in the case of Ramsey policies – providing a positive definition is more of a challenge. Most settings of interest will be more complex than the inflation bias example above, and will not admit entirely constant policies. Endogenous state variables, in particular, will generally evolve over time, changing the set of possible choices. The question: 'symmetry along what dimension?' is therefore likely to arise.

Symmetry will be particularly relevant to the treatment of forward-looking constraints over time. In any given period t, changes to current and future policy can affect the realised values of the functions in constraint (5) in current, future *and past* periods. Under what circumstances could it be meaningful to assert that these constraints are being affected by 'the same amount' by a given change?

Drawing on well-known approaches from the literature on comparability and social choice, we extend the definition of the h function to specify the extent to which its values can be compared across states and time, and the extent to which its values just capture normalisations. The basic idea is that it will not make sense to seek a 'symmetric' solution where the definition of symmetry rests on an arbitrary normalisation of h. More positively, if there are comparisons that *can* be asserted to have economic meaning, it may be possible to exploit these to construct a notion of dynamic symmetry in the treatment of forward-looking constraints.

3.4.1 Mathematical equivalence

Mathematically, it is always possible to find a general set of transformations to the h and h^0 functions, together with the intertemporal discount factor, that preserve inequality (5) in all states and initial time periods. Though the main h function that we have defined is time-invariant, these transformations could in principle vary according to the period in which the function in question is evaluated, and the state σ. So long as the set of underlying allocations that satisfy the condition is unaffected, the basic choice problem will not change.

Formally, let the function $\phi(\cdot; t, \sigma) : \mathbb{R} \to \mathbb{R}$ be interpreted as a transformation to the specific h function that features in period t and state σ. Likewise, let $\phi^0(\cdot; s, \sigma') : \mathbb{R} \to \mathbb{R}$ be interpreted as a transformation to h^0 for the constraint (5) that features in period s for state σ'.

If this pair of transformation functions is to leave the mathematical structure of the constraint set unaffected, the following inequality must be equivalent to condition (5):

$$\begin{aligned} & \mathbb{E}_t \left[\sum_{\tau=0}^{\infty} \tilde{\beta}_{t,t+\tau}\left(\sigma, \sigma^{+\tau}\right) \phi\left(h\left(a_{t+\tau}\left(\sigma^{+\tau}\right), \sigma^{+\tau}\right), t, \sigma^{+\tau}\right) \middle| \sigma \right] \\ \geq \quad & \phi^0\left(h^0\left(a_t\left(\sigma\right), \sigma\right), t, \sigma\right) \end{aligned} \tag{7}$$

where $\tilde{\beta}_{t,t+\tau}(\sigma, \sigma^{+\tau}) \in \mathbb{R}$ is a transformation to the discount factor applying between periods t and $t+\tau$, given the relevant states. An allocation will satisfy (7) if and only if it satisfies (5). Recall that condition (5) states:

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

A simple example of a transformation with this property would be that generated by arbitrary scalar addition to h:

$$\begin{aligned}\phi\left(h,t,\sigma^{+\tau}\right) &= h+\alpha_{t+\tau}\left(\sigma^{+\tau}\right)\\ \phi^{0}\left(h^{0},t,\sigma\right) &= h^{0}+\mathbb{E}_{t}\left[\sum_{\tau=0}^{\infty}\beta^{\tau}\alpha_{t+\tau}\left(\sigma^{+\tau}\right)\middle|\sigma\right]\\ \tilde{\beta}_{t,t+\tau}\left(\sigma,\sigma^{+\tau}\right) &= \beta^{\tau}\end{aligned}$$

where $\alpha_t(\sigma)\in\mathbb{R}$ for all t and σ, subject to the discounted expected sum being bounded.

Another possibility is multiplication by a time-varying coefficient:

$$\begin{aligned}\phi\left(h,t,\sigma^{+\tau}\right) &= \delta_{t+\tau}h\\ \phi^{0}\left(h^{0},t,\sigma\right) &= \delta_{t}h^{0}\\ \tilde{\beta}_{t,t+\tau}\left(\sigma,\sigma^{+\tau}\right) &= \beta^{\tau}\frac{\delta_{t}}{\delta_{t+\tau}}\end{aligned}$$

where $\delta_t\in\mathbb{R}$.

3.4.2 Economic equivalence

A **comparability assumption** will be a requirement for policy choice to be unaffected by a subset of the possible ϕ transformations, labelled 'admissible' transformations. The subset of admissible transformations will be more restrictive than the complete set that leaves (5) unaffected, and its exact specification will be context-specific.

There are two main comparability possibilities that we will consider. The first, and most widely applicable, is difference comparability.

Definition. The function h is **difference comparable** if:

$$h\left(a_t'\left(\sigma\right),\sigma\right)-h\left(a_t''\left(\sigma\right),\sigma\right)=h\left(a_s^{*}\left(\sigma'\right),\sigma'\right)-h\left(a_s^{**}\left(\sigma'\right),\sigma'\right)$$

implies:

$$\phi\left(h\left(a_t'\left(\sigma\right),\sigma\right),t,\sigma\right)-\phi\left(h\left(a_t''\left(\sigma\right),\sigma\right),t,\sigma\right)=\phi\left(h\left(a_s^{*}\left(\sigma'\right),\sigma'\right),s,\sigma'\right)-\phi\left(h\left(a_s^{**}\left(\sigma'\right),\sigma'\right),s,\sigma'\right)$$

for all admissible ϕ functions, given any pair of periods t and s, states $\{\sigma,\sigma'\}\in\Sigma^2$, and policies $\{a_t'(\sigma),a_t''(\sigma),a_s^{*}(\sigma'),a_s^{*}(\sigma')\}\in A_\sigma^4$.

Thus difference comparability implies that the levels of any changes to the h functions have meaning relative to one another, when compared across time and states. It is meaningful to describe one increase in h as greater or lower in magnitude than another. In classic social choice problems, difference comparability across agents' utility functions is necessary to justify utilitarian or weighted utilitarian social welfare criteria. 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 utility functions is required for the social welfare function to be meaningful, and if it is assumed for the functions that feature in social objective then it makes sense to extend it to the (identical) functions that feature in the constraints.

Difference comparability is also an appropriate assumption to make when treating linearised models. If the h function included in the model represents the first-order deviation of a non-linear function from some reference level, changes to that reference level should not affect policy, at least so long as the underlying linear structure of the model remains valid. This is guaranteed by difference comparability. Thus, for instance, it implies that models expressed in terms of an output gap would not be affected by a choice to express the gap relative to an efficient equilibrium level versus a flexible-price equilibrium level, holding parameters constant.

If the h function is difference comparable, it is defined up to the class of affine transformations of the form:

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

where the scalar $\delta \in \mathbb{R}$ is common across time and states, but the additive coefficient $\alpha_t\left(\sigma\right) \in \mathbb{R}$ can vary in both. Corresponding to this are the following ϕ^0 function and discount factor $\tilde{\beta}_{t,t+\tau}\left(\sigma, \sigma^{+\tau}\right)$:

$$\begin{aligned}
\phi^0\left(h^0, t, \sigma\right) &= \delta h^0 + \mathbb{E}_t\left[\sum_{\tau=0}^{\infty} \beta^\tau \alpha_{t+\tau}\left(\sigma^{+\tau}\right) \middle| \sigma\right] \\
\tilde{\beta}_{t,t+\tau}\left(\sigma, \sigma^{+\tau}\right) &= \beta^\tau
\end{aligned}$$

It is straightforward to show that no other transformations are consistent with preserving difference comparability.

An alternative to difference comparability is for h to be ratio comparable.

Definition. The function h **is ratio comparable** if:

$$\frac{h\left(a_t'\left(\sigma\right), \sigma\right)}{h\left(a_t''\left(\sigma\right), \sigma\right)} = \Delta$$

for some and $\Delta \in \mathbb{R}$ implies:

$$\frac{\phi\left(h\left(a_t'\left(\sigma\right), \sigma\right), t, \sigma\right)}{\phi\left(h\left(a_t''\left(\sigma\right), \sigma\right), t, \sigma\right)} = \Delta$$

for all admissible ϕ functions, given any period t, state σ and policies $\left\{a_t'\left(\sigma\right), a_t''\left(\sigma\right)\right\} \in A_\sigma^2$.

This assumption is appropriate for h functions where proportional increases are clearly defined, even if the units in which h is measured may not be. Notice that it requires a 'zero' value for h to have independent economic meaning. A change to the intercept will always affect ratio comparisons, so this must be ruled out of the class of admissible transforms.

An example of an h function with defined ratios is one that specifies an agent's net expenditure from assets 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. Another example is a firm's within-period profit function.

For h to have defined ratios, the class of admissible ϕ functions must be restricted to those of the form:

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

where $\delta_t(\sigma) \in \mathbb{R}$ for all t and σ. The corresponding ϕ^0 function and discount factor are then given by:

$$\begin{aligned} \phi^0\left(h^0, t, \sigma\right) &= \delta_t(\sigma) h^0 \\ \tilde{\beta}_{t,t+\tau}\left(\sigma, \sigma^+\right) &= \beta^\tau \frac{\delta_t(\sigma)}{\delta_{t+\tau}(\sigma^{+\tau})} \end{aligned}$$

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

3.5 The feasible set

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

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

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

It is also convenient to distinguish separately the set of allocations that is consistent with the conventional, structural constraints in (5), and the set of allocations that is consistent with the forward-looking constraints (5) alone. The set of allocations from period s onwards that satisfy the forward-looking constraints (5) only will be denoted Ξ^h:

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

Note that this set is independent of any initial state vector, since by assumption these do not feature in constraint (5). Similarly, the set of allocations that satisfy constraints (4) only for any given x_{s-1} will be denoted $\Xi^g(x_{s-1})$:

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

Note that $\Xi(x_{s-1})$ is the intersection space between these two constraint sets:

$$\Xi(x_{s-1}) = \Xi^h \cap \Xi^g(x_{s-1})$$

3.6 Feasibility and possibility

In any given environment there will usually be a possibility to eliminate certain variables and constraints mathematically, so that the dimensionality of the problem is reduced. Analytically, this amounts to replacing some of the constraints that define $\Xi(x_{s-1})$ with restrictions on the underlying space of possible allocations $(\mathcal{X} \times \mathcal{A})$. For most purposes the choice between 'constraint' and 'basic space' amounts to an innocuous normalisation, but it can make a difference in the event that some of the *infeasible* allocations in $(\mathcal{X} \times \mathcal{A})$ are invoked in the analysis. This is true, in particular, if choice is required to be independent of irrelevant alternatives (IIA), which 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. *(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 $\varrho^h : \Sigma \to R^h(\sigma)$, $\varrho^{h^0} : \Sigma \to R^{h^0}(\sigma)$ and $\varrho^r : \Sigma \to R^r(\sigma)$ 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 feasibility constraints.

We stress that this normalisation is sufficient, but by no means necessary for our purposes. In every example that we have studied it is possible to impose extensive cross-restrictions on the universe of possible values for h, h^0 and r, and still invoke IIA to the full extent required. The assumption provides the most general guarantee possible that *non-existence* of irrelevant alternatives will never impede the analysis. Less expansive definitions of A will serve equally well on a case-by-case basis.

3.7 Structural assumptions

To place structure on the problem, we will work with differing combinations of the following assumptions on the main primitives:

Assumption 2. *The functions r, p, 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 3. *p is quasi-concave.*

Assumption 4. *h is concave and h^0 is convex.*

Assumption 5. *r is strictly concave.*

Assumption 2 provides essential structure and is imposed throughout. In most environments of interest the relevant r, p, h, and h^0 functions are utility functions, production functions, profit functions and the like, for which continuity is a conventional assumption. Convexity of the basic spaces A_σ and X is similarly uncontentious. Compactness of A_σ and X is a stronger assumption, as it 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 3 ensures that the constraint set described by the 'standard' feasibility restriction (3) is convex. This will be useful for deriving sufficiency statements in particular. Assumptions 4 and 5 are quite strong: they require concavity, not quasiconcavity. These are likewise needed to obtain some important sufficiency results, though necessary conditions for time-consistently undominated policy will be possible without them. As ever, a failure to satisfy such strong requirements in a particular example does not mean that our analysis is irrelevant to it – just that some conditions may require more work for confirmation.[23]

Finally, Assumption 5 imposes strict concavity on the r function whereas Assumption 4 only places weak concavity/convexity on h and h^0. All results would survive the alternative of weak concavity in r and strict concavity/convexity in h and h^0.

4 Desirable sets

The focus of our paper is on normative choice in the general setting presented above. Clearly this is complicated by the time inconsistency problem, and making progress requires some formalisation of this problem, and the options that it leaves open.

Time consistency as a concept can be defined by reference to the set of optimal choices. Denote by $C(x_{s-1}) \subseteq \Xi(x_{s-1})$ the 'argmax' set of policies, optimal from the perspective of s, given x_{s-1}. That is:

$$C(x_{s-1}) = \arg \max_{(\bar{\mathbf{x}}_\mathbf{s}, \bar{\mathbf{a}}_\mathbf{s}) \in \Xi(x_{s-1})} W_s$$

An equivalent definition of $C(x_{s-1})$ can be offered in terms of the binary relation $\succeq$, and this provides a useful point of departure for the arguments that follow:

$$C(x_{s-1}) = \{(\bar{\mathbf{x}}_\mathbf{s}, \bar{\mathbf{a}}_\mathbf{s}) \in \Xi(x_{s-1}) : \forall (\bar{\mathbf{x}}'_\mathbf{s}, \bar{\mathbf{a}}'_\mathbf{s}) \in \Xi(x_{s-1}), (\bar{\mathbf{x}}_\mathbf{s}, \bar{\mathbf{a}}_\mathbf{s}) \succeq (\bar{\mathbf{x}}'_\mathbf{s}, \bar{\mathbf{a}}'_\mathbf{s})\}$$

That is, $C(x_{s-1})$ is the set of allocations in $\Xi(x_{s-1})$ that weakly dominate all other feasible allocations according to $\succeq$.

[23]This is closely connected to the well-known difficulty of establishing sufficient conditions for maxima in optimal tax settings, where the constraint set is nonconvex. Following Lucas and Stokey (1983), analysis usually proceeds by focusing on the implications of necessary optimality conditions, checking computed examples for alternative maxima numerically. An equivalent approach to this is possible here.

Time consistency obtains whenever $(\bar{\mathbf{x}}^*_{\mathbf{s}}, \bar{\mathbf{a}}^*_{\mathbf{s}}) \in C(x_{s-1})$ implies$(\bar{\mathbf{x}}^*_{\mathbf{t}}, \bar{\mathbf{a}}^*_{\mathbf{t}}) \in C\left(x^*_{t-1}\right)$ for all $t > s$. Thus a best policy today is a best policy to continue with tomorrow. Note that time consistency does *not* require that the ordering $\succeq$ should remain stable through time. Even in time-consistent environments, it is quite possible that two allocations $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ in $\mathcal{X} \times \mathcal{A}$ will be such that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \succeq (\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ but $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}}) \succ (\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ for $t > s$.[24]

As is well known, sufficient conditions for time consistency to obtain are that (a) social preferences are recursive, as we assume throughout, and (b) all of the feasible continuation allocations available in period $t > s$ should also have been available from the perspective of period s, given the allocation chosen between s and $t-1$.

Kydland and Prescott problems are environments where the second of these conditions fails, because some continuation allocations that are feasible in t will be inconsistent with actions taken prior to t. This means that recursive membership of the argmax set will not generally be possible. *Ex-ante* it is desirable to make promises that rule out certain future choices; *ex-post* it is not optimal to keep these promises.

Given this, there are two options for normative choice. The first is to define an 'initial' period s, and commit to an allocation in $C(x_{s-1})$, regardless of its subsequent desirability properties. Since *recursive* membership of the argmax set is not possible, this option settles for membership of argmax in the initial period instead. This is the conventional, Ramsey approach to designing optimal policy, and the properties of the associated allocations have been explored extensively across a number of different policy literatures.

A second option, which has not been pursued in the literature to date, is to try to expand the 'desirable' set of allocations in any given period, beyond $C(x_{s-1})$. Recursive membership of a sufficiently expanded set may be possible where recursive membership of the argmax set C is not. Heuristically, if no policy is recursively 'best', perhaps there may nonetheless exist policies that are recursively 'satisfactory', somehow defined. This is the general normative approach that we explore in the current paper.

Formally, we will proceed to construct a set $D(x_{s-1}) \supseteq C(x_{s-1})$ for all s, interpreted as a more weakly desirable set than $C(x_{s-1})$. The main vehicle for doing this is a binary relation $\succeq^{TC}$ that we define on the space of dynamic allocations in Ξ^h – i.e., the set of allocations consistent with the model's forward-looking constraints. The relation $\succeq^{TC}$ will be incomplete, applying only between pairs of allocations that, we assert, can be ranked in a 'time-consistent' manner. The basic idea is to isolate comparisons across allocations that are uncontentious, at least with respect to the time inconsistency problem. If an allocation is to be chosen in period s, then there should not be an alternative allocation available in s that dominates the chosen one according to the strict

[24]Heuristically, this would be the case if the relative benefits from $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ were relatively 'frontloaded'.

relation $\succ^{TC}$. Thus, given this incomplete ordering, $D(x_{s-1})$ will satisfy:[25]

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

As Section 5 details, there will be two conditions under which $\succeq^{TC}$ is asserted. The first is when comparing allocations that deliver equivalent values to one another for the constraint functions that feature in the forward-looking restriction (5). If choice is restricted to a subset of allocations with this property, then by construction there is no possibility that future decisions will violate past promises. Time inconsistency does not arise, and so choice can be made in a standard way. The relation $\succeq^{TC}$ will coincide with $\succeq$ across all pairs in these sub-domains.

The second possibility arises when a pair of allocations is ranked in the same way at every point in time according to the standard social preference ordering $\succeq$. For instance, suppose that the two allocations $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ and $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ are both feasible (in continuation) for every $t \geq s$, and $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}}) \succ (\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ always holds. In this case there is complete unanimity over time in the comparison between the two allocations, and no preference-based justification exists for selecting $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ over $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$.[26] We will assert that $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}}) \succ^{TC} (\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ applies in this case.

As noted, a necessary condition for an allocation belonging to the set $D(x_{s-1})$ is that no other feasible choice should dominate it under the ordering $\succeq^{TC}$. By itself this has the weakness that more restrictive feasibility conditions can limit the scope for dominance to arise, meaning that the 'desirable' set $D(x_{s-1})$ can grow in size simply by ruling out seemingly irrelevant options. To avoid this, we impose an 'independence of irrelevant alternatives' (IIA) requirement directly on $D(x_{s-1})$. This amounts to defining an expanded set of options $\hat{\Xi}(x_{s-1}) \supseteq \Xi(x_{s-1})$ that includes all additional choices in Ξ^h that are dominated by options in $\Xi(x_{s-1})$ according to the ranking $\succ^{TC}$. These options, though dominated, may themselves dominate alternative options in $\Xi(x_{s-1})$. Given this, $D(x_{s-1})$ will be fully characterised by:

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

In words, a feasible allocation belongs to the desirable set if it is not dominated by an alternative feasible option, or by an irrelevant infeasible option.

To summarise, our response to the impossibility of finding a policy that is recursively *optimal* will be to seek policies that are recursively *undominated*, according to an incomplete relation $\succeq^{TC}$. This realtion will be defined in a way that, we argue, captures the set of time-consistent policy comparisons that can be made in spite of general time inconsistency. This approach contrasts with Ramsey policymaking, which seeks a policy that is optimal in the model's initial

[25]The reason for the inequality $\subseteq$ here is that an IIA restriction is additionally used to limit $D(x_{s-1})$, as explained below.

[26]As explained below, in this case the social preference ranking can be considered 'timeless', with neither the feasible set nor social preferences varying as time progresses.

period alone, and need not satisfy any desirability criterion *per se* thereafter. Clearly the appeal of our method will stand or fall on the appeal of the criteria we use to define $\succeq^{TC}$, and it is to this that we now turn.

5 Time-consistent orderings

This section explains the situations in which the ordering $\succeq^{TC}$ will apply.

5.1 Basic definition and rationality properties

The relation $\succeq^{TC}$ will be defined on selected pairs in the restricted space $\Xi^h \subseteq (\mathcal{X} \times \mathcal{A})$. Note that this implies it need not be stable if the forward-looking constraints in the problem vary, but it is invariant to the feasibility restrictions that make up Ξ^g. The former is to be expected: the time inconsistency problem is central to the need to define time-consistent dominance in the first place. $\succeq^{TC}$ will be reflexive, but need not be complete on any given $\Xi(x_{s-1})$. Since our focus will be on undominated allocations, transitivity is unnecessary and not imposed.[27] Strict and indifference orderings $\succ^{TC}$ and $\sim^{TC}$ are also defined, and relate to $\succeq^{TC}$ in the usual way.

Given a feasible set $\Xi^g(x_{s-1})$, $\succeq^{TC}$ will characterise a set of undominated allocations $D(x_{s-1})$ as above.

5.2 Time-consistent comparisons

Time consistency obtains when an optimal selection from a set of feasible allocations is stable over time. The construction of the ordering $\succeq^{TC}$ is based on the idea that even if choice from the entire feasible set is not time-consistent, choice between particular sets of options may have this property.

To preview the approach in more detail, we will assert two distinct possibilities for ordering pairs of allocations with $\succeq^{TC}$. The first possibility is that the two allocations can be compared without any time inconsistency problem arising. Loosely, this happens when a pair of allocations imply equivalent values for the constraint functions h and h^0. This means that switching between the two policies will continue to respect any past commitments, whatever these may have been. In this situation standard choice procedures generate no inconsistency, and so we can impose $\succeq^{TC}$ if and only if the original preference ordering $\succeq$ applies. This a 'constraint-based' approach to defining dominance – based on finding restricted constraint sets where time inconsistency does not apply.

The second approach is to compare allocations based on consistent preference rankings over time. Suppose that there is a pair of allocations $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ with the property that $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ is feasible in period t if and only if

[27] That is, for an allocation to be dominated via transitivity, it must already be dominated without it. A subtler issue arises if feasibility rules out the original dominating allocation, but this will be addressed via a formal IIA condition placed on $D(x_{s-1})$, rather than by using transitivity on $\succeq^{TC}$.

$(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ is likewise. Thus conditional on one of the two allocations being chosen, there is meaningful option to switch to the other at each point in time. If one of these allocations is strictly preferred to the other for all $t \geq s$, we will assert that $\succ^{TC}$ applies. By contrast with the first, constraint-based approach, this provides a 'preference-based' route to time consistency. By construction, the argmax set in the restricted choice between $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ and $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ will only ever contain the continuation of the allocation that was considered optimal between $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ in period s.

The next two subsections formalise these approaches in turn.

5.3 Constraint-based comparisons

5.3.1 Promises

An integral component of the analysis that follows will be the link between an allocation and a sequence of **promises**. Promise values have long been a central analytical device when considering time inconsistency problems, whether in the recursive characterisation of Ramsey choices or the analysis of 'sustainable plans' without a commitment device. We will also make use of them, but in doing so will retain the focus on sequence-space analysis. This contrasts with the usual recursive formulation of problems that use promise values.

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

$$\omega'_t(\sigma_-) := \mathbb{E}_{t-1}\left[\sum_{\tau=0}^{\infty} \beta^{\tau} h\left(a'_{t+\tau}\left(\sigma^{+\tau}\right), \sigma^{+\tau}\right) \middle| \sigma_-\right] \tag{8}$$

where $\sigma^{+\tau}$ denotes a successor history to σ_-, $\tau+1$ periods ahead. The promise value is defined as the expected discounted value of the h function from period t onwards, assessed in the prior period $t-1$, given some history σ_- up to $t-1$. The sequence of values $\{\omega'_t(\sigma)\}_{t=s}^{\infty}$ will belong to the set of bounded sequences l^{∞} for any given $\sigma \in \Sigma$, and $\mathcal{W}$ denotes the product space of these sequences across σ. Since the function h is continuous and A_{σ} is compact, a uniform upper bound for $\omega_t(\sigma)$ will exist, independent of σ.

If the allocation $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ is also in the set Ξ^h, it must satisfy the forward-looking constraint (5) for all date-states. Rewriting this in terms of promises gives:

$$h(a'_t(\sigma), \sigma) + \beta\omega'_{t+1}(\sigma) \geq h^0(x'_t, a'_t(\sigma), \sigma) \tag{9}$$

for all $\sigma \in \Sigma$ and all t. In addition, note that the definition of the promises (8) can be rewritten in a recursive form:

$$\omega'_t(\sigma) = \mathbb{E}_{t-1}\left[h(a'_t(\sigma'), \sigma') + \beta\omega'_{t+1}(\sigma') \middle| \sigma\right] \tag{10}$$

where σ' is used to denote a successor history to σ.

5.3.2 Equivalent representations

As explained in Section 3.4, there will generally be a set of admissible transformations of the h function that preserve its economic interpretation, and to which policy choice should be invariant. Corresponding to these will be a sequence of admissible transformations to promise values, which we denote $\phi^{\omega}(\omega, t, \sigma)$. Specifically, if an admissible transformation ϕ is applied to the h function, $\omega_t(\sigma)$ must be transformed to $\phi^{\omega}(\omega_t(\sigma), t, \sigma)$ in a corresponding manner in order to preserve consistency.

Specifically, suppose that the h function is **levels comparable**. Then it is immediate that the set of admissible transforms to $\omega_t(\sigma)$ will take the form:

$$\phi^{\omega}(\omega_t(\sigma), t, \sigma) = \delta\omega_t(\sigma) + \alpha_t(\sigma)$$

where $\alpha_t(\sigma) \in \mathbb{R}$, and $\delta \in \mathbb{R}$ is common across shocks and time. This implies that the sequence of differences between any pair of promise sequences, $\{\omega'_t(\sigma) - \omega''_t(\sigma)\}_{t=s}^{\infty}$, is defined up to the scalar δ. Other comparisons, such as sequences of ratios, are not defined. Technically this will be particularly helpful when analysing differential changes to a promise sequence, since it provides a well-defined vector movement away from any given sequence $\{\omega_t(\sigma)\}_{t=s}^{\infty}$, independent of the chosen $\alpha_t(\sigma)$ used.

Now suppose instead that the h function is **ratio comparable**. In this case the class of admissible transforms to $\omega_t(\sigma)$ take the form:

$$\phi^{\omega}(\omega_t(\sigma), t, \sigma) = \delta_t(\sigma)\,\omega_t(\sigma)$$

where $\delta_t(\sigma) \in \mathbb{R}$. In this case the sequence of ratios of promises, $\left\{\frac{\omega'_t(\sigma)}{\omega''_t(\sigma)}\right\}_{t=s}^{\infty}$ is defined irrespective of any admissible transformation. Technically this means that a sequence of differential changes to the promises will have meaning if expressed as a proportional increment over existing values.

5.3.3 Time-consistent comparisons

Promises play a useful normative role because they characterise subsets of the constraint set in which no time inconsistency problem arises. Consider an allocation $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in \Xi^h$, inducing promises $\bar{\omega}'_{\mathbf{s}}$. Suppose that there is another allocation $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$, satisfying the following two restrictions for all $t \geq s$ and all $\sigma \in \Sigma$:[28]

$$h(x''_t, a''_t(\sigma), \sigma) + \beta\omega'_{t+1}(\sigma) \geq h^0(x''_t, a''_t(\sigma), \sigma) \tag{11}$$
$$\mathbb{E}_{t-1}\left[h(x''_t, a''_t(\sigma'), \sigma') + \beta\omega'_{t+1}(\sigma') \middle| \sigma\right] \geq \omega'_t(\sigma) \tag{12}$$

Condition (11) states that the promises implied by the allocation $(\bar{\mathbf{x}}'_{\mathbf{t+1}}, \bar{\mathbf{a}}'_{\mathbf{t+1}})$ are consistent with satisfying the forward-looking constraint in period t when the within-period choices for that period are(x''_t, a''_t). This means that switching

[28]Note that these restrictions directly imply that $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \in \Xi^h$.

from the allocation (x''_t, a''_t) in period t to $\left(\bar{\mathbf{x}}'_{\mathbf{t+1}}, \bar{\mathbf{a}}'_{\mathbf{t+1}}\right)$ after period t would not violate the forward-looking constraints that applied in t. Condition 12 states that the within-period allocation (x''_t, a''_t) implements values for the h functions that are at least as great as the allocation (x'_t, a'_t). This means that switching the within-period allocation in t from (x'_t, a'_t) to (x''_t, a''_t) would not imply that any forward-looking constraints prior to t had been violated.

Taken together, the two restrictions imply that a free choice in each period $t \geq s$ between the within-period allocations (x'_t, a'_t) and (x''_t, a''_t) would guarantee satisfaction of all forward-looking constraints that could possibly be generated by these two allocations from s onwards. In addition, by (12), the resulting choice would be guaranteed to satisfy any forward-looking restrictions that applied *prior* to s, provided these are satisfied by $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$.

This motivates the following definition of tome-consistent comparability between allocations:

Definition. The allocation $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ is **time-consistently comparable** to the allocation $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in \Xi^h$ if conditions (11) and (12) hold for all $t \geq s$ and all $\sigma \in \Sigma$.

Note that this definition is not symmetric – that is, $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ may be time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ without the converse holding. This is because the sequence $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ could imply a 'tougher' set of promise values $\bar{\omega}''_{\mathbf{s}}$, such that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ does not satisfy condition 12 with these promises. In this case some forward-looking restrictions prior to s might be satisfied by $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ but not by $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$. The purpose of defining comparability is to characterise the set of alternative allocations that could potentially be found to dominate $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$. Since the dominance relation will itself be asymmetric, this initial asymmetry does not imply any inconsistencies.

Another feature of the definition is that time-consistent comparability is itself consistent through time. That is, if $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$, then $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ for all $t > s$. This follows mechanically from the definition, as conditions (11) and (12) are required to hold for all $t \geq s$, but it is an essential property for ensuring time-consistent choice across alternative sequences.

As their name suggests, time-consistently comparable allocations are useful for constructing restricted choice environments where the time inconsistency problem does not apply. To explain this fully, it helps to define composite allocations as follows:

Definition. The allocation $(\bar{\mathbf{x}}^*_{\mathbf{s}}, \bar{\mathbf{a}}^*_{\mathbf{s}})$ is a **composite** of $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ if for all $t \geq s$,$(x^*_t, a^*_t) \in \{(x'_t, a'_t), (x''_t, a''_t)\}$.

That is, a composite allocation is constructed by taking the within-period choice from one or other of these two sequences in each period t. Figure 10 provides an abstract visualisation. The allocation $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ is charted in green and $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ in red. A composite allocation is some arbitrary path between the two. One example is mapped out by the arrows.

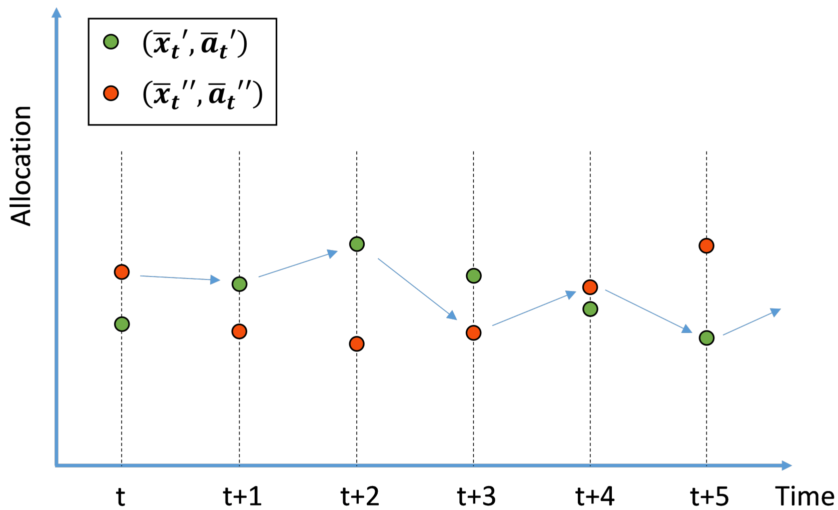

Figure 10: A composite allocation

The following observation follows immediately from the definition of time-consistent comparability:

Remark. Suppose $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in \Xi^h$. Then every composite of $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ is also time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$, and belongs to Ξ^h.

Being able to extend the 'time-consistent comparability' status immediately to composite allocations is extremely useful, because the availability of composites is central to proving time-consistent choice. Time consistency means that the optimal selection from a set of possible dynamic allocations does not depend on the time period in which that choice is assessed. Suppose that choice were allowed only between two distinct allocations $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$, but not their composites. Then it is quite possible that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \succ (\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ in period s, but $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}}) \succ (\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ in period $t > s$. Time consistency would not be assured. Once all composites are allowed, however, an optimal continuation policy from t onwards can – and indeed must – always form part of an optimal plan from the perspective of period s. Proposition 2 makes direct use of this argument.

To this end, let $\mathcal{T}_s \subseteq \Xi^h$ be a set of allocations from period s onwards, with $\mathcal{T}_t \subseteq \Xi^h$ denoting the corresponding set of continuations from $t > s$ on. Completeness of $\mathcal{T}_s$ is defined as follows:

Definition. $\mathcal{T}_s$ is **complete** if for every pair $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}), (\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \in \mathcal{T}_s$, every composite of $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ also belongs to $\mathcal{T}_s$.

A characteristic feature of the Kydland and Prescott problem is that the set Ξ^h is *not* complete. Allocations from t onwards must be consistent with forward-looking constraints prior to t, and this cross-restriction rules out switches from some dynamic paths to others. Time-consistent comparability is a useful concept precisely because it generates subsets of Ξ^h that *are* complete. As noted above, if $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in \Xi^h$, then every composite of the two allocations belongs to Ξ^h (and is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$). Choice on the resulting sub-domains will thus be time-consistent.

To prove this point formally, it helps to extend the definition of time-consistent comparability to cover sets. Thus the set $\mathcal{T}_s$ will be described as time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ if $\mathcal{T}_s$ contains $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$, and every element of $\mathcal{T}_s$ is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ as defined above.

We then have the following result. The proof is very simple, and illustrates the equivalence between completeness and time-consistency. Hence it is left in the main text.

Proposition 2. *Fix $x_{s-1} \in X$ and $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in \Xi(x_{s-1})$. For any complete set of allocations $\mathcal{T}_s$ that is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$, if $(\bar{\mathbf{x}}^*_{\mathbf{s}}, \bar{\mathbf{a}}^*_{\mathbf{s}}) \in \arg\max_{\mathcal{T}_s \cap \Xi^g(x_{s-1})} W_s$ then $(\bar{\mathbf{x}}^*_{\mathbf{t}}, \bar{\mathbf{a}}^*_{\mathbf{t}}) \in \arg\max_{\mathcal{T}_t \cap \Xi^g\left(x^*_{t-1}\right)} W_t$ for all $t > s$.*

Proof. Suppose there were an alternative allocation $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ that came to dominate $(\bar{\mathbf{x}}^*_{\mathbf{t}}, \bar{\mathbf{a}}^*_{\mathbf{t}})$ by period $t > s$. Then the composite allocation $\left(\left(\bar{\mathbf{x}}^{\mathbf{t-1}*}_{\mathbf{r}}, \bar{\mathbf{x}}'_{\mathbf{t}}\right), \left(\bar{\mathbf{a}}^{\mathbf{t-1}*}_{\mathbf{r}}, \bar{\mathbf{a}}'_{\mathbf{t}}\right)\right)$ would be strictly superior to $(\bar{\mathbf{x}}^*_{\mathbf{s}}, \bar{\mathbf{a}}^*_{\mathbf{s}})$ 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 $(\bar{\mathbf{x}}^*_{\mathbf{s}}, \bar{\mathbf{a}}^*_{\mathbf{s}})$ being in the arg max set. $\square$

Thus complete sets of time-consistently comparable allocations are domains in which time-consistent choice obtains.

5.3.4 Imposing an ordering

Having isolated sub-domains in which time inconsistency does not apply, it remains to assert that the ordering $\succeq^{TC}$ will coincide with the time-consistent ranking $\succeq$ within these sub-domains. Formally:

Condition 1. (Constraint dominance) For every $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in \Xi^h$, if $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \in \Xi^h$ is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ then $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \succ (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ implies $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \succ^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$, and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \sim (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ implies $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \sim^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$.

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

Our concern throughout this paper is the Kydland and Prescott time inconsistency problem. If this problem does not affect the comparison among a restricted set of options, we have no reason to depart from a standard selection procedure within that restricted domain. Condition 1 ensures that we do not. It implies that if one policy implements a given set of promises at a strictly lower social welfare cost than another, the latter will be dominated according to the ordering $\succ^{TC}$. If both are feasible, the dominated option will not belong to $D(x_{s-1})$, and so will not be selected.

A separate question is whether this 'standard selection procedure' is normatively appropriate in general, independently of time inconsistency. A long tradition dating back to Ramsey (1928) has argued that social choice ought to give greater weight to the welfare of future generations than sequential optimisation with discounting would permit.[29] Though it is easy to have sympathy with this perspective, its implications stretch far wider than the set of Kydland and Prescott problems. One of the main reasons for proceeding as we do is to highlight the possibilities for treating the time inconsistency problem distinctly from the wider question of appropriate social discounting. This further motivates using standard selection criteria whenever time inconsistency is absent.

5.3.5 Implication: an 'inner problem'

An implication of Condition 1 is that the following problem becomes central to the analysis:

Problem 1. (Inner Problem)

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

subject to:

$$\begin{aligned} \mathbb{E}_{t-1}\left[h\left(a_t'\left(\sigma'\right), \sigma'\right) + \beta\omega_{t+1}\left(\sigma'\right) \middle| \sigma\right] &\geq \omega\left(\sigma\right) && (13) \\ h\left(a_t\left(\sigma\right), \sigma\right) + \beta\omega_{t+1}\left(\sigma\right) &\geq h^0\left(a_t\left(\sigma\right), \sigma\right) && (14) \end{aligned}$$

for all $t \geq s$ and all $\sigma \in \Sigma$, with $\bar{\omega}_{\mathbf{s}} \in \mathcal{W}$ and $x_{s-1} \in X$ given.

The main interest in Problem 1 comes from the following, two-part Proposition.

Proposition 3. *1. For any $x_{s-1} \in X$, suppose that the ordering $\succ^{TC}$ satisfies Condition 1. Then each allocation in $D(x_{s-1})$ solves Problem 1, given the promise values that it induces.*

2. Fix a state vector $x_{s-1} \in X$, and suppose the allocation $(\bar{\mathbf{x}}_{\mathbf{s}}', \bar{\mathbf{a}}_{\mathbf{s}}')$ solves Problem 1, given the promise values that it induces. Then $(\bar{\mathbf{x}}_{\mathbf{s}}', \bar{\mathbf{a}}_{\mathbf{s}}') \in D(x_{s-1})$ is consistent with Condition 1.

Proposition 3 is divided into complementary 'necessity' and 'sufficiency' parts. Part 1 states that if Condition 1 is applying, solving Problem 1 for the promise values that it induces is a necessary feature of every allocation in $D(x_{s-1})$. This is immediate. Every allocation in the constraint set for Problem 1 is time-consistently comparable with an allocation that induces the given promise values. So if an allocation does not solve Problem 1 for the promise values it induces, it must be constraint-dominated – and hence does not belong to $D(x_{s-1})$.

[29] Recently this issue has figured large in the policy literature on climate change (Stern, 2006), and the dynamic social insurance literature, where discounting has been identified as an important source of the immiseration result (see Phelan, 2006, and Farhi and Werning, 2007).

Part 2 states that *if* an allocation solves Problem 1 for the promise values that it induces, then membership of $D(x_{s-1})$ by this allocation is consistent with Condition 1. Intuitively, this works because the constraint set for Problem 1 includes *all* allocations that could conceivably dominate one that induces the chosen promise values.

Overall, the Proposition can be read as implying that 'undominated under Condition 1' and 'solving Problem 1' are near-interchangeable requirements. There can be no allocation in $D(x_{s-1})$ that fails to solve Problem 1 when Condition 1 applies. Conversely, solving Problem 1 for its own promise values is enough to guarantee that an allocation is not dominated under Condition 1. For practical purposes the latter of these will be the more useful. It will be easier to confirm that an allocation solves Problem 1 for the promise values that it induces than to try to check for membership of $D(x_{s-1})$ more directly.

The generality of this result is qualified by the requirement that an allocation should solve Problem 1 *for the promise values that it induces*. It remains possible that a sequence $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$, inducing promises $\bar{\omega}'_{\mathbf{s}}$, could solve Problem 1 for promises $\bar{\omega}''_{\mathbf{s}} \neq \bar{\omega}'_{\mathbf{s}}$, but not for $\bar{\omega}'_{\mathbf{s}}$.[30] In this case 'solving Problem 1' and 'undominated under Condition 1' are not equivalent. A sufficient condition for this not to matter is that constraint (13) should be binding for all σ and all t, since this directly implies that the solution implies the given promise values. In practical applications it will be straightforward to check that the corresponding multipliers are indeed positive.

Problem 1 is one component in a two-part division of the overall choice problem. It is referred to in what follows as the **inner problem**, because it is concerned with the optimal choice of allocation *given* a sequence of promise values. By design, it is entirely time-consistent, and could be solved by recursive optimisation. The **outer problem** is the problem of choosing a sequence of promise values, $\bar{\omega}_{\mathbf{s}}$.

5.4 Preference-based comparisons

5.4.1 Motivation: timeless choice domains

The idea behind constraint dominance was to find subsets of Ξ^h in which conventional time-consistent choice was bound to obtain, given recursive preferences. This could be done provided the subsets were complete, so that an optimal continuation from t onwards must always have been available from the perspective of a prior period s. This motivated Condition 1, which asserts that standard time-consistent choice techniques should be observed where the constraint set thus allows.

In this subsection we identify an alternative route to time-consistent choice, based directly on information from preference rankings. Within the set of allocations Ξ^h, there may exist sub-groups in which social preference orderings

[30] This happens because a slack inequality in (13) in period $t+1$ allows scope for inequality (14) to be relaxed in t: an increase in $\omega_{t+1}(\sigma)$ will not rule out the optimal choice. But this will include new allocations in the constraint set, some of which may be superior.

are completely unchanging through time. A simple example comes from the linear-quadratic inflation bias problem of section 2. Consider there the set of allocations in Ξ^h that hold inflation and output *constant* in all periods. If in period s a permanent inflation-output combination of (π', y') is preferred to an alternative (π'', y''), this ranking will remain the same in period t.

Now suppose that choice were being analysed just across a subset of options where this form of dynamic agreement is present, and with the additional restriction that all allocations within the subset are available at all points in time. In such circumstances the time dimension effectively becomes irrelevant to the problem. There is a unique, conventional preference ordering across some indexed space of options. There may be interesting economic dynamics associated with some or all of the choices available, but these are just descriptive features of the different possibilities. If consumer preferences do not depend on the weather, there is no need to include rainfall as a variable in consumer choice analysis. The idea here is to isolate environments where the same is true of time.

Whenever a subset of options has this property, the ordering $\succeq^{TC}$ will be asserted to coincide with the time-invariant social preference ranking. The qualification that all allocations must be available at all points in time is important in the context of state variables, since these have the potential to change feasibility restrictions as time progresses.

As the inflation bias example should indicate, the emphasis here is not on finding choice sets from which the optimal choice would necessarily follow in equilibrium from recursive decisionmaking. If the current inflation rate were chosen under the assumption that future policymakers will have equivalent freedom, the specification of choice across complete dynamic sequences does not make sense. In the example, the current policymaker could not do anything to guarantee constant future outcomes.[31] Instead, the focus is on domains where 'once-and-for-all' decisions, if possible, would never once be regretted. This may be viewed as a form of time consistency, but perhaps a better term for such choice settings is 'timeless'.

5.4.2 Time-invariant feasibility

Formally the purpose is to extend the binary relation $\succeq^{TC}$ beyond the range of comparisons implied by Condition 1 alone. Recall that the domain for this relation is Ξ^h, the set of allocations consistent with the forward-looking restrictions (5).

A complication in isolating timeless choice comparisons is the possibility that feasibility may vary over time, due to the presence of state variables. In period s, given some initial x_{s-1}, it may be that both $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ are feasible paths. In period $t > s$, given x'_{t-1} $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ will remain feasible by definition, but $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ may not be. This is possible whenever $x'_{t-1} \neq x''_{t-1}$. Whether or not preferences are stable, in this case choice between $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ certainly

[31] Indeed, if future re-optimisation were anticipated, sub-optimal discretionary equilibria would presumably follow in the usual way.

depends on time. In s there is a meaningful choice between the two allocations, and in t there is not.

We restrict the definition of dominance to comparisons where a meaningful pairwise choice is guaranteed to exist in all periods, provided it exists in the first. To do this without direct reference to the feasibility constraints requires that the relation $\succeq^{TC}$ should be defined conditional on the sequence of state variables that a pair of allocations induces. Formally:

Definition. Allocations $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ in Ξ^h are said to be **timelessly comparable** in period s if $\bar{\mathbf{x}}'_{\mathbf{s}} = \bar{\mathbf{x}}''_{\mathbf{s}}$.

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

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

The feasibility of a pair of timelessly 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 $t-1$.[32]

5.4.3 Preference dominance defined

The dominance relation $\succeq^{TC}$ will be asserted whenever preferences between a pair of timelessly comparable allocations are completely unchanging over time. Again, the justification is that the time index has effectively become redundant in describing the choice problem across these allocations. Normative choice within this restricted domain is unambiguous, so long as current and future social preferences are assumed to be all that matter to it.

There is one technical complication that emerges in assessing whether preferences between a pair of allocations are unchanging through time. Given the infinite-horizon setting, it may be possible for strict preference between a pair of allocations to hold in every *finite* time period, but indifference at the limit as $t \to \infty$. The implication is that the marginal benefits from switching away from the inferior allocation fade to zero as time progresses, under any welfare metric. This is clearly a boundary case between time-consistent strict preference and time-varying preference, and we opt not to assert any ordering when it arises. This restricts the applicability of $\succeq^{TC}$, and so potentially allows more allocations to be included in the undominated set $D(x_{s-1})$ in any period s.

The main analytical implication of this choice is that time-invariant strict preference must be defined by reference to boundedness from the lower contour sets of different allternatives. Let $\mathcal{L}(\bar{\mathbf{a}}_{\mathbf{s}}; \bar{\mathbf{x}}'_{\mathbf{s}}) := \{\bar{\mathbf{a}}'_{\mathbf{s}} \in \mathcal{A} : (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}_{\mathbf{s}}) \succeq (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})\}$ be the lower contour set for the allocation $\bar{\mathbf{a}}_{\mathbf{s}}$, holding constant the sequence of

[32] Clearly if a different allocation entirely is pursued from s to $t-1$, it may be that neither $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ nor $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ belongs to $\Xi(x_{t-1})$ for the corresponding x_{t-1}. The definition provides alternative allocations that it is always possible to switch among, but off-path feasibility can never be assured.

state vectors at $\bar{\mathbf{x}}'_{\mathbf{s}}$. If the norm on $\mathcal{A}$ is denoted by $\|\cdot\|$, then from the definition of a lower contour set we have that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \succ (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ applies if and only if there exists an $\varepsilon > 0$ such that $\|(\bar{\mathbf{a}}''_{\mathbf{s}} - \bar{\mathbf{a}}_{\mathbf{s}})\| \geq \varepsilon$ for all $\bar{\mathbf{a}}_{\mathbf{s}} \in \mathcal{L}(\bar{\mathbf{a}}'_{\mathbf{s}}; \bar{\mathbf{x}}'_{\mathbf{s}})$. This can be extended to ensure time-invariant strict preference, including at the limit, by asserting a *uniform* ε bound from the lower contour set over time.

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

Condition 2. (Preference dominance) For any $x_{s-1} \in X$ and pair of timelessly comparable allocations $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}), (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \in \Xi^h$:

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

2. If $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}}) \sim (\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ for all $t \geq s$, then $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \sim^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$.

Thus if $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ is strictly preferred to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ at every point in time, $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \succ^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ will apply. If the two allocations are viewed with indifference at each point in time, $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \sim^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ will apply. In all other cases the condition is silent. If $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \succ^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ holds by application of Condition 2, we say that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ **preference-dominates** $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$.

With either option feasible at each point in time, implementing a dominated allocation from period s onwards would represent a clear failure of normative choice. Qualitatively, this is precisely the sort of failure that is well-known to arise when policy is set period-by-period under discretion. An unambiguously better option always exists, and is never chosen. A minimum requirement for normative choice is that it should improve on this coordination failure.

5.4.4 Constraint dominance versus preference dominance

Notice that there is a symmetry between the domains in which Conditions 1 and 2 can be applied. Constraint dominance allows comparisons between allocations that implement the same sequence of promises over time. Preference dominance allows comparisons between allocations that implement the same sequence of states through time. Since promises are often treated as auxiliary states, this symmetry has superficial appeal.

Within their respective domains, however, constraint dominance has far wider applicability. *Any* pair of sequences that implement the same promises will be ranked according to Condition 1, whereas Condition 2 only applies when there is also complete agreement through time in social preferences.

6 Undominated choices and promises

This section explores the link between the set of undominated policies $D(x_{s-1})$ and the choice problem across alternative promise sequences $\bar{\omega}_{\mathbf{s}}$.

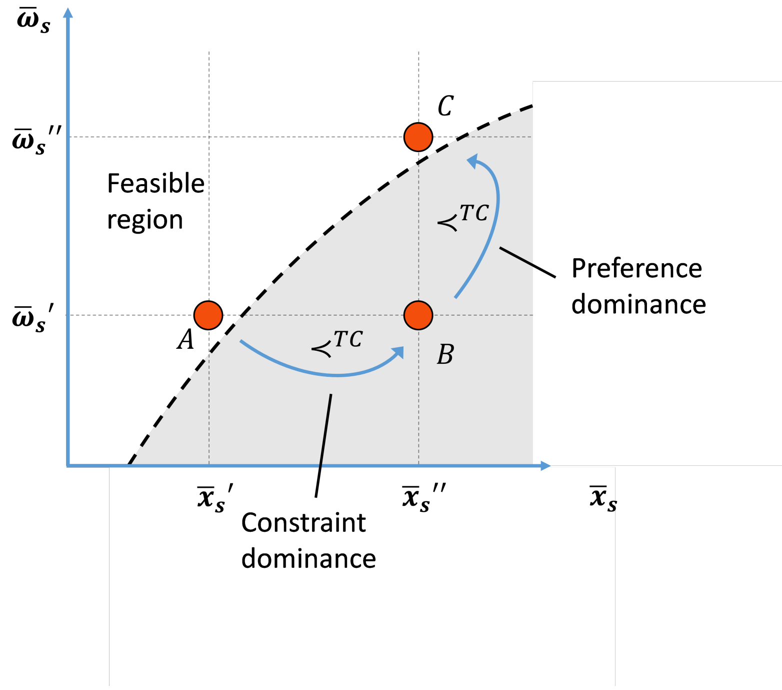

Figure 11: Pairwise comparisons by alternative criteria

6.1 Independence of irrelevant alternatives

The ordering $\succeq^{TC}$ provides the basis for constructing a set of undominated policies $D(x_{s-1})$ for any $x_{s-1} \in X$. This is based on two conditions – constraint dominance and preference dominance – that are quite restrictive in their applicability, and in some settings fesibility may severely limit the set of comparisons that can actually be made. The issue is highlighted in Figure 11. This shows three allocations, A, B, and C, such that $B \succ^{TC} A$ by constraint dominance, and $C \succ^{TC} B$ by preference dominance. The axes capture the fact that constraint dominance requires the same promise values to be applied by the two allocations, and preference dominance requires the same evolution of the state vector. A feasibility frontier is given by the dashed line, with allocations A and C feasible and B not feasible.

Restricting attention – for heuristic purposes – to these three allocations alone, if B is not feasible then neither allocation A nor allocation C is dominated by an alternative feasible choice, and the set $D(x_{s-1})$ would contain both A and C. But if the set of feasible allocations were expanded to include B, A would no longer belong to undominated set. This is true even though the additional alternative B is not itself a member of the undominated set. The implication is that $D(x_{s-1})$ as defined may be sensitive to the inclusion of 'irrelevant' alternatives, particularly in environments where it is not feasible to move promises without moving states, or vice-versa.

Independence of choice from the inculsion of irrelevant alternatives is a conventional rationality property, and so this example motivates a refinement of the set $D(x_{s-1})$ to rule out options that would be dominated if feasibility were expanded to include other dominated choices. To this end, we use the following definition:

Definition. Fix $x'_{s-1} \in X$, and let $\Xi^g_c\left(x'_{s-1}\right)$ be the complement of $\Xi^g\left(x'_{s-1}\right)$ in $\mathcal{X} \times \mathcal{A}$. The set $\tilde{\Xi}^g\left(x'_{s-1}\right) \supset \Xi^g\left(x'_{s-1}\right)$ is an **irrelevant extension** of $\Xi^g\left(x'_{s-1}\right)$ if for every $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in \tilde{\Xi}^g\left(x'_{s-1}\right) \cap \Xi^g_c\left(x'_{s-1}\right) \cap \Xi^h$, for all $t \geq s$ there is an allocation $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}}) \in \Xi^g\left(x'_{t-1}\right) \cap \Xi^h$ such that $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}}) \succ^{TC} (\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$.

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 existing allocation, and thus none of the new allocations belongs to the set of undominated allocations for the expanded constraint set. Note that there is no requirement for the dominating allocation $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ to be the same through time, i.e. $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ need not be the period-t continuation of the equivalent $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$.

This allows a more restrictive definition of the set of undominated allocations, limiting attention only to those that survive the inculsion of irrelevant alternatives in the feasible set. The desirable set $D(x_{s-1})$ can then be defined formally as follows:

Condition 3. Fix $x_{s-1} \in X$. $D(x_{s-1})$ is the set of allocations in $\Xi^g(x_{s-1}) \cap \Xi^h$ that are undominated according to $\succ^{TC}$ in every set $\tilde{\Xi}^g(x_{s-1}) \cap \Xi^h$ such that $\tilde{\Xi}^g(x_{s-1})$ is an irrelevant extension of $\Xi^g(x_{s-1})$.

This condition is distinct in form from 1 and 2, as it is used directly to characterise the desirable set rather than to establish the underlying ordering $\succ^{TC}$. Indeed, it is dependent on this ordering having already been defined. There is a symmetry in the analytical approach, however. Conditions 1 and 2 are motivated by the observation that choice in some subsets of $\Xi(x_{s-1})$ ought to be uncontroversial. This was either because a time inconsistency problem did not apply to these subsets, or because preferences within them were completely unchanging through time. Extending the set of possibilities beyond these subsets may introduce new, difficult choices, but it should not change the rankings that are already established.

Condition 3 applies a similar independence logic in reverse. If a selection can be made from a large set of options, shrinking the options by eliminating unchosen alternatives should not change choice. For Conditions 1 and 2 the idea is that the *in*clusion of extra options should not change existing rankings; here it is that the *ex*clusion of dominated alternatives should not affect choice.

6.2 An equivalent condition on promises

Proposition 3 demonstrated the tight connection between choosing an undominated policy according to Condition 1, and solving for an optimal allocation given a complete sequence of promises. Viewed through this dual lens, adding Conditions 2 and 3 allows for something additionally to be said about the appropriate choice of the promise sequence.

The policymaker's preference ordering over alternative allocations $\succeq$ induces a preference ordering over the space of possible promise sequences, under the assumption that for any promise sequence that is selected, Problem 1 will be solved to determine an allocation. Formally, we have the following definition:

Definition. Fix $x_{s-1} \in X$, and let $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ respectively solve Problem 1 when the promise sequences are $\bar{\omega}'_{\mathbf{s}}$ and $\bar{\omega}''_{\mathbf{s}}$ and the lagged state vector is x_{s-1}. The ordering $\succeq^{\omega}_{x_{s-1}}$ is defined by:

$$\bar{\omega}'_{\mathbf{s}} \succeq^{\omega}_{x_{s-1}} \bar{\omega}''_{\mathbf{s}} \longleftrightarrow (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \succeq (\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$$

Note that the solution to Problem 1 depends on the initial state vector x_{s-1}, and so the relative ranking of $\bar{\omega}'_{\mathbf{s}}$ and $\bar{\omega}''_{\mathbf{s}}$ may thus, in turn, depend on x_{s-1}. The subscript on the relation $\succeq^{\omega}_{x_{s-1}}$ emphasises this dependence. Since it is derived from the complete preference relation $\succeq$, $\succeq^{\omega}_{x_{s-1}}$ will be complete on the set of promise sequences for which Problem 1 has a solution.

Given the ordering $\succeq^{\omega}_{x_{s-1}}$, a lower contour set $\mathcal{L}^{\omega}$ can be defined in promise space as follows:

$$\mathcal{L}^{\omega}\left(\bar{\omega}'_{\mathbf{s}}; x_{s-1}\right) := \left\{\bar{\omega}_{\mathbf{s}} \in \mathcal{W} : \bar{\omega}'_{\mathbf{s}} \succeq^{\omega}_{x_{s-1}} \bar{\omega}_{\mathbf{s}}\right\}$$

A parallel requirement to Condition 2 can then be placed on the choice of the promise sequence $\bar{\omega}_{\mathbf{s}}$, based on dynamic agreement in the ordering $\succeq^{\omega}_{x_{s-1}}$.

Definition. Fix $x'_{s-1} \in X$, and consider the promise sequence $\bar{\omega}'_{\mathbf{s}}$ such that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ solves Problem 1 for x'_{s-1} and $\bar{\omega}'_{\mathbf{s}}$. $\bar{\omega}'_{\mathbf{s}}$ is **time-consistently dominated** by the alternative $\bar{\omega}''_{\mathbf{s}}$ if and only if there exists an $\varepsilon > 0$ such that for all $t \geq s$, $\|\bar{\omega}''_{\mathbf{t}} - \bar{\omega}_{\mathbf{t}}\| \geq \varepsilon$ for all $\bar{\omega}_{\mathbf{t}} \in \mathcal{L}^{\omega}\left(\bar{\omega}'_{\mathbf{t}}; x'_{t-1}\right)$.

The interest in this notion of dominance derives from the following Proposition:

Proposition 4. *Fix $x_{s-1} \in X$, and let $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ solve Problem 1 for the promise sequence that it induces, denoted $\bar{\omega}'_{\mathbf{s}}$. When the ordering $\succ^{TC}$ satisfies Conditions 1 and 2, the allocation $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ belongs to $D\left(x_{s-1}\right)$ if and only if $\bar{\omega}'_{\mathbf{s}}$ is not time-consistently dominated by any alternative promise sequence.*

This Proposition develops the link between the choice of undominated policies under $\succ^{TC}$, and the choice of promises for Problem 1. Proposition 3 showed that constraint dominance was essentailly equivalent to requiring that Problem 1 should be solved for *some* promise sequence. Proposition 4 shows how preference dominance and the independence of irrelevant alternatives refinement allow something further to be said about the *choice* of promise sequence.

Specifically, if it is possible to find a promise sequence that is not time-consistently dominated by any alternative, this is almost enough to ensure that the allocation solving Problem 1 for this promise sequence is not dominated according to $\succ^{TC}$. The only qualification is that the solution must induce the promise sequence in question, which is equivalent to saying that the promise constraints should (weakly) bind. Since promise constraints can be tightened until this is true, this qualification is not a significant impediment – a point that confirmed more formally by Proposition 12 below.

Though there are clear parallels between the definition of preference dominance for allocations and the definition of time-consistent dominance for promise sequences, the two are not the same – and thus the equivalence is not trivial. In particular, suppose the promise sequence $\bar{\omega}'_{\mathbf{s}}$ is time-consistently dominated by $\bar{\omega}''_{\mathbf{s}}$, given some initial x'_{s-1}, and let $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ solve Problem 1 for $\bar{\omega}'_{\mathbf{s}}$ and $\bar{\omega}''_{\mathbf{s}}$ respectively. The fact that $\bar{\omega}'_{\mathbf{s}}$ is time-consistently dominated by $\bar{\omega}''_{\mathbf{s}}$ does not place any restriction on the evolution of the state vector $\bar{\mathbf{x}}''_{\mathbf{s}}$, and thus it is not the case that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ need necessarily preference-dominate $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$. Relatedly, the allocation $\left(\tilde{\bar{\mathbf{x}}}''_{\mathbf{t}}, \tilde{\bar{\mathbf{a}}}''_{\mathbf{t}}\right)$ that solves Problem 1 for $\bar{\omega}''_{\mathbf{t}}$ and x'_{t-1} must be superior to the contiuation $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$, but there is no requirement in the definition that $\left(\tilde{\bar{\mathbf{x}}}''_{\mathbf{t}}, \tilde{\bar{\mathbf{a}}}''_{\mathbf{t}}\right)$ should be the period-t contiuation of $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$. If x''_{t-1} has departed from x'_{t-1}, it is quite possible that the period-t continuation of $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ will not be feasible in t.

There is a close link between the equivalence results here and envelope conditions. When $\succ^{TC}$ embeds the preference dominance condition, membership of D^* means that it must not be possible to vary the chosen allocation in the subspace $\mathcal{A}$ alone and improve social welfare at every horizon. When an allocation solves Problem 1 for a given promise sequence, the time path for the state variables is chosen optimally, and varying the chosen promise sequence will generally imply variation in the chosen sequence for the states. But for small changes in the promises, any marginal gains that are attainable by changing promises should remain attainable when the sequence for state variables is required to remain constant. Thus the absence of time-consistent gains from varying promises corresponds to the absence of preference dominance when states are held fixed. The fact that promise-based comparisons allow states to be varied does not allow extra scope for improvement at the margin, because states are already chosen optimally. The next section develops a value function representation of the problem that allows this intuition to be formalised.

7 A value function representation

Propositions 3 and 4 together suggest that the problem of choosing undominated allocations can be divided into an 'inner' problem of choosing optimal allocations for given promise sequences, and an 'outer' problem of choosing promise sequences that are not time-consistently dominated. The inner problem subsumes the time-consistent aspects of choice, and can be analysed by standard techniques. The outer problem is less conventional. This section develops some apparatus that will be helpful in characterising it. We focus on two specific objects: the feasible set of promise sequences, and the value function associated with a choice of promise sequence. This subsection defines these objects, and characterises the necessary restrictions on the economic environment in order for them to satisfy useful regularity properties.

7.1 Feasible sets of promise sequences

When analysing the outer problem it is useful to be able to refer to a **feasible set** of promise sequences – those that could potentially be selected from any period $s \geq 0$ onwards. For some choices of $\bar{\omega}_{\mathbf{s}}$ the constraint set for Problem 1 may be empty – there simply does not exist a feasible policy that can make good on these promises. Clearly such promise sequences are not feasible selections.

Formally, we denote the set of feasible promise sequences from s onwards by $\Omega\left(x_{s-1}\right)$. This is defined as follows for any given $x_{s-1} \in X$:

$$\Omega\left(x_{s-1}\right) := \left\{\bar{\omega}_{\mathbf{s}} \in \mathcal{W} : \text{ constraint set to Problem 1 nonempty}\right\}$$

For ideas of continuity and differential changes to promises, it is useful to refer to the interior of $\Omega\left(x_{s-1}\right)$. This is denoted by $\mathring{\Omega}\left(x_{s-1}\right)$. Though possible in principle, we have not encountered examples where the choice of a desirable promise value is constrained by the boundary of $\Omega\left(x_{s-1}\right)$. Intuitively, a promise sequence at the boundary of $\Omega\left(x_{s-1}\right)$ is one for which the constraint set of Problem 1 is arbitrarily close to being empty, meaning that the burden of keeping promises is almost too great to be sustained by available resources. In most examples this will not be a desirable calibration of promises. For this reason we will generally focus on results that relate to promise choices in $\mathring{\Omega}\left(x_{s-1}\right)$.

7.1.1 Convexity of $\Omega\left(x_{s-1}\right)$

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

Proposition 5. *Suppose Assumptions 2, 4 and 3 hold. Fix $x_s \in X$. The space $\Omega\left(x_{s-1}\right)$ 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 6 below. Note that assumption 4 is relatively strong, requiring concavity in the constraint function h rather than quasiconcavity, and convexity in the case of h^0.

7.2 Value of the inner problem

The maximised **value** of the inner problem is denoted by $V(\bar{\omega}_{\mathbf{s}}; x_{s-1})$, for all $\bar{\omega}_{\mathbf{s}} \in \Omega\left(x_{s-1}\right)$ and all $x_{s-1} \in X$. For all $\omega_{\mathbf{s}} \in \mathcal{W}$ not in $\Omega\left(x_{s-1}\right)$, we normalise $V\left(\bar{\omega}_{\mathbf{s}}; x_{s-1}\right)$ to -∞ for convenience. Note that V can be viewed as a cardinalisation of the preference ordering $\succeq^{\omega}_{x_{s-1}}$, 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.

7.2.1 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 6. *Suppose Assumptions 2, 4, 5 and 3 hold. Fix $x_{s-1} \in X$. $V(\cdot; x_{s-1})$ is* ***concave*** *in $\bar{\omega}_{\mathbf{s}} \in \Omega(x_{s-1})$.*

As in Proposition 5, which concerned the convexity of $\Omega(x_{s-1})$, placing this additional structure on V does not come without a cost. Assumption 4 implies concavity in the h function and convexity in the h^0 function, rather than quasi-concavity and quasi-convexity respectively – something that is not always guaranteed in the initial representation of the problem. As for Proposition 7, renormalising the h (h^0) function to endow it with concavity (convexity) would allow the mathematical problem to be overcome, but this may come with a loss of economic interpretability. As ever, when the required assumptions are not satisfied the analysis can proceed, but with caveats. The most direct analogy 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 1. *Suppose the assumptions for Proposition 6 are true. Fix $x_{s-1} \in X$, and let $\bar{\omega}'_{\mathbf{s}}$ and $\bar{\omega}''_{\mathbf{s}}$ be arbitrary selections from $\Omega(x_{s-1})$. Then $V(\alpha\bar{\omega}'_{\mathbf{s}} + (1-\alpha)\bar{\omega}''_{\mathbf{s}}; x_{s-1})$ 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.

7.2.2 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 $\bar{\omega}_{\mathbf{s}} \in \Omega(x_{s-1})$, the directional (Gateaux) derivative of V is denoted by $\delta_V(\bar{\omega}_{\mathbf{s}}, x_{s-1}; \bar{\mathbf{w}}_{\mathbf{s}})$, defined for all $\bar{\mathbf{w}}_{\mathbf{s}} \in \mathcal{W}$ by:[33]

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

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

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

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. As shown by Marcet and Marimon (1998, 2016), in many settings of interest these multipliers have an intuitive connection to Pareto weights in a cross-sectional allocation problem. It is particularly informative in our context to contrast the evolution of these Pareto weights over time under alternative procedures for choosing $\bar{\omega}_{\mathbf{s}}$.

In general we denote the present-value multiplier on promise-keeping constraint (13) for history σ in period t by $\lambda_t^k(\sigma)$, and the corresponding promise-making constraint (14) 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 $\bar{\lambda}_{\mathbf{s}}^{\mathbf{k}}$ and $\bar{\lambda}_{\mathbf{s}}^{\mathbf{m}}$ are infinite sequences of these from s on. With some abuse of notation, the space that $\bar{\lambda}_{\mathbf{s}}^{\mathbf{k}}$ and $\bar{\lambda}_{\mathbf{s}}^{\mathbf{m}}$ inhabit is denoted $\mathcal{W}^*$.[34]

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.[35] Formally, we will make use of the following:

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

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

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 $\bar{\omega}_{\mathbf{s}}$ lies at the boundary of $\Omega(x_{s-1})$.[36] In addition, the condition rules out the simple incorporation of equality constraints as two-sided inequalities, since in this case interiority is impossible. Straightforward extensions to the main arguments are possible that allow for linear forward-looking constraints,[37] but we neglect these to avoid over-complicating the analysis.[38]

Proposition 7. *Suppose Assumptions 1, 2, 4, 5 and 3 hold. Fix $x_{s-1} \in X$, and let $\bar{\omega}_{\mathbf{s}} \in \Omega(x_{s-1})$ be such that the constraint set for Problem 1 contains an inner point. Then wherever the directional derivative $\delta_V(\bar{\omega}_{\mathbf{s}}, x_{s-1}; \bar{\mathbf{w}}_{\mathbf{s}})$ exists*

[34] Strictly this does not correspond to the dual of $\mathcal{W}$, because the multipliers are expressed in present-value form.

[35] See, for instance, Luenberger (1969),§8.3, Theorem 1.

[36] If 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.

[37] Since concavity in the h function and convexity in the h^0 function are needed for the results below, two-sided inequalities will only work under the assumption of linearity.

[38] Linear 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.

and is finite-valued, there is a pair of Lagrange multiplier sequences $\bar{\lambda}^{\mathbf{k}}_{\mathbf{s}}$ *and* $\bar{\lambda}^{\mathbf{m}}_{\mathbf{s}}$ *in* $\mathcal{W}^*$ *such that* $\delta_V(\bar{\omega}_{\mathbf{s}}, x_{s-1}; \bar{\mathbf{w}}_{\mathbf{s}})$ *is given by:*

$$\delta_V(\bar{\omega}_{\mathbf{s}}, x_{s-1}; \bar{\mathbf{w}}_{\mathbf{s}}) = \sum_{t=s}^{\infty} \beta^{t-s} \int_{\sigma\in\Sigma} \left\{\beta\left[\lambda^m_t(\sigma) + \lambda^k_t(\sigma_-)\right] w_{t+1}(\sigma) - \lambda^k_t(\sigma) w_t(\sigma)\right\} d\Pi(\sigma) \tag{15}$$

with σ_- *the predecessor history to* σ*.*

All of the major characterisation results that follow will assume differentiability in V, applying condition (15), but a generalisation to points of non-differentiability would be technically straightforward. Where the derivative $\delta_V(\bar{\omega}_{\mathbf{s}}, x_{s-1}; \bar{\mathbf{w}}_{\mathbf{s}})$ does not exist, there is a set of Lagrange multipliers $\mathbf{\Lambda}^{\mathbf{k}}_{\mathbf{s}} \times \mathbf{\Lambda}^{\mathbf{m}}_{\mathbf{s}} \subset \mathcal{W}^* \times \mathcal{W}^*$ such that $\delta^+_V(\bar{\omega}_{\mathbf{s}}, x_{s-1}; \bar{\mathbf{w}}_{\mathbf{s}})$ is the minimum in $\mathbf{\Lambda}^{\mathbf{k}}_{\mathbf{s}} \times \mathbf{\Lambda}^{\mathbf{m}}_{\mathbf{s}}$ of the object on the right-hand side of (15), and $\delta^-_V(\bar{\omega}_{\mathbf{s}}, x_{s-1}; \bar{\mathbf{w}}_{\mathbf{s}})$ is its maximum.

Condition (15) is ultimately a straightforward 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 $\beta\left[\lambda^m_t(\sigma) + \lambda^k_t(\sigma_-)\right]$.[39] Against this, an increase in $\omega_t(\sigma)$ tightens the promise-keeping restriction in period t. This accounts for the term $\lambda^k_t(\sigma)$.

8 Policy characterisation

This section considers necessary and sufficient properties for policies that inhabit the set $D(x_{t-1})$ for all time periods, and contrasts these with the properties of Ramsey policy.

Specifically, the results show that time-consistently undominated policy can be characterised by a long-run relationship between the multipliers on the promise-making and promise-keeping constraints, (14) and (13). Under the restrictions that we have used to define time-consistent dominance there is no unique characterisation of the transition to any steady state, and multiple transition paths will generically exist even when steady state is unique. This multiplicity is addressed directly in Section 9.4. But the long-run characterisation is sufficient to demonstrate that Ramsey policy will not generally satisfy the requirements for being time-consistently undominated. Since Ramsey policy must belong to the undominated set $D(x_{s-1})$ when it is first devised, the implication is that its longer-term dynamics are inconsistent with it remaining undominated. This has important consequences for the use of long-run policy conclusions that arise from solving the Ramsey problem, including classic zero capital tax results. The allocation that is implied by Ramsey policy in steady state is time-consistently dominated by alternative options, and thus cannot be considered desirable *per se*.

[39] Recall that $\omega_{t+1}(\sigma)$ enters into the promise-keeping constraint (13) for the predecessor state σ_- in t, whereas it enters the promise-making constraint (14) for σ itself.

8.1 Ramsey policy

Before proceeding it is helpful to re-cast the Ramsey problem as a choice over alternative promise sequences. Given x_{s-1}, this problem solves:

$$\max_{\bar{\omega}_{\mathbf{s}} \in \Omega(x_{s-1})} V\left(\bar{\omega}_{\mathbf{s}}; x_{s-1}\right)$$

Equation (15) above characterises the derivatives of V with respect to the promise values. Assuming the optimal choice is interior, and that the relevant multipliers are uniquely defined, the solution to this problem is characterised by the usual requirement for V to be flat along every feasible dimension of movement for the promises. Heuristically, when σ has positive measure a first-order condition can be recovered by setting $w_t(\sigma)$ to 1 in (15) for a given σ and t, and zero elsewhere. Consistent with this, Ramsey policy will generally require:

$$\lambda_t^k(\sigma) = \lambda_{t-1}^m(\sigma) + \lambda_{t-1}^k(\sigma_-) \tag{16}$$

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

$$\lambda_s^k(\sigma) = 0 \tag{17}$$

for Π-almost all $\sigma \in \Sigma$.

Condition (16) is familiar from the influential work of Marcet and Marimon (1998, 2016), who build on this multiplier characterisation to show how Ramsey-optimal policy can be recovered from a recursive saddle-point problem, using the objects on the right-hand side of (16) to augment the state vector. The time inconsistency of the solution is apparent. By (16) the promise-keeping multipliers must be non-decreasing in t, but in the initial period s they are set to zero. Any re-optimisation subsequent to s must imply a different choice of promise sequence, so that the equivalent of (17) is satisfied. Intuitively, it is never desirable *ex-post* to pay a positive shadow cost to keep prior promises.

Time inconsistency of this kind is well understood. A more important consideration for the current paper is whether 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, we have the following Proposition:

Proposition 8. *Let $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ 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 $t > s$, $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}}) \notin D\left(x'_{t-1}\right)$.*

The intuition behind the result relates directly to the fundamental time inconsistency problem. When a Ramsey policy is chosen in period s, the promises that are selected for all periods after s equalise the *ex-ante* benefits from making a promise with the *ex-post* costs of keeping them. If the chosen sequence of promises is reassessed in any subsequent period $t > s$, this trade-off remains optimally struck for the future, but not for t itself. In period t there is a costly promise being kept, with no remaining benefits.

It follows that all policymakers from $s+1$ onwards wish to reduce the burden of the promises that they are required to keep, and they may be willing to do this even if it comes at the cost of some sub-optimality in future promise choice. Indeed, at the margin, the cost of altering *future* promises is negligible, relative to the Ramsey-optimal choice, precisely because future promises have been chosen optimally. Thus a Pareto improvement will be possible, whereby each period's policymaker accepts that future promises will be somewhat less stringent, in return for an easing of present constraints. The possibility of this improvement rules out that $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ could belong to $D\left(x'_{t-1}\right)$.

The only clear exception to this is a situation in which the promise multipliers approach zero under Ramsey policy – either in finite time or at the limit. In this case there is no strict Pareto improvement that can be obtained, because future policymakers are not constrained by the proimises they keep. But this is a case in which the time inconsistency problem has ceased to matter, so it shouldn't be surprising that optimal choice appears less problematic.

8.2 Time-consistently undominated policy: necessity

The most significant of the general 'necessity' results that we provide for time-consistently undominated policy choices is the following:

Proposition 9. *Suppose that the policy $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ is time-consistently undominated, given some initial $x'_{s-1} \in X$, and assume that V is differentiable at the induced promise sequence $\bar{\omega}'_{\mathbf{s}}$. If h is* ***difference-comparable*** *then for Π-almost all $\sigma \in \Sigma$, either:*

1. *There is no period τ such that both $\left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right]$ 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:*

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

for all $r \geq T$.

If instead h is ***ratio-comparable*** *then for Π-almost all $\sigma \in \Sigma$, either:*

1. *There is no period τ such that both $\left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right]\omega_{t+1}(\sigma)$ and $\lambda_t^k(\sigma)\omega_t(\sigma)$ 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:*

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

for all $r \geq T$.

We focus the discussion principally on the difference-comparable case. Most elements remain relevant with ratio comparability. With difference comparability, the Proposition should be read as a statement about the long-run tendency of the ratio $\frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}$. So long as this ratio exists and is bounded above zero in every period from some τ onwards, its compounded product from τ onwards must be stable relative to any non-trivial geometric process. More informally, this can be interpreted as implying a long-run gravitation in the multipliers towards the relationship:

$$\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]=\lambda_t^k(\sigma) \tag{18}$$

Specifically, there must either be long-run convergence to (18) holding with equality, or each of the inequalities $\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right] \gtrless \lambda_t^k(\sigma)$ must hold an infinite number of times, so that the dynamic product of the ratio $\frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}$ remains stable.

In general there is nothing to guarantee that the multiplier values need necessarily converge to a steady state under time-consistently undominated policy, though in practice convergent cases will be more straightforward to analyse. If they do converge then the ratio $\frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}$ must converge in value to unity, but limit cycles – or more exotic dynamics – around a value of one for this object are equally consistent with the result.[40]

Analytically, the result is obtained by taking differential changes to the promise sequences, and showing that these can deliver strict improvements for all current and future policymakers when the boundedness condition is not met. The requirement that both $\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]$ 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.

There is an important separation in condition (18) between the state histories σ and σ_-, and calendar time t. It is commonly argued that the benefits from policy commitment in Kydland and Prescott environments derive from the ability to ensure 'history dependence' in choice – i.e., dependence of policy in period t on what happened prior to t. The Ramsey condition (16) guarantees this, through the dependence of current multipliers on their realised past values. At first glance there would seem to be an inherent tension between choosing a policy that is time-consistently desirable (in any sense) and providing this history dependence. Time consistency implies that a choice in t should be justifiable solely by reference to outcomes from t onwards, whereas history dependence requires dynamic linkages. But as condition (18) indicates, it is possible to link promise multipliers between successive exogenous histories within a given

[40] There is a clear analogue in the possibility of limit cycles and chaotic dynamics in Pareto-efficient equilibria for standard two-period OLG models, as highlighted by Grandmont (1985).

calendar time period. The marginal cost of keeping a promise for exogenous history σ in t can be linked to the marginal benefit from making an equivalent promise for that history in period t, to be realised in $t+1$. There is nothing *optimal* about linking marginal costs and benefits across periods in this way, but the results above show that doing so may at least allow a time-consistently undominated choice.

When the h functions are ratio comparable rather than difference comparable, the main change to the boundedness condition is an allowance for the possibility of a change to the units in which $\omega_t(\sigma')$ is expressed over time. Suppose , for instance, that $\omega_t(\sigma')$ represents a consumer's wealth level, and there is inflation in the price index over time. This creates a secular tendency in the promise values to increase in numerical value over time, and in the multipliers to decrease.[41] The condition implies boundedness in the multiplier ratio once this has been corrected for.

In the event that convergence in the multipliers does occur, a far sharper statement is possible, anticipated by condition (18):[42]

Corollary 2. *Suppose that the policy sequence $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ is time-consistently undominated, and induces multipliers $\lambda^k_t(\sigma)$ and $\lambda^m_t(\sigma)$ that converge to bounded steady-state values $\lambda^k_{ss}(\sigma)$ and $\lambda^m_{ss}(\sigma)$ for all $\sigma \in \Sigma$. Then:*

$$\beta\left[\lambda^m_{ss}(\sigma) + \lambda^k_{ss}(\sigma_-)\right] = \lambda^k_{ss}(\sigma) \tag{19}$$

As stated this Corollary is restricted to the case of difference comparability in the h function. Under ratio comparability the same applies, provided there is additionally convergence to a steady-state value in the promise values. The main purpose in stating the Corollary is to highlight that the more mathematically involved form of Proposition 9 really flows from its focus on general promise sequences, with no restriction that convergence need occur. When convergence is assured, time-consistently undominated policy mandates a very simple equality restriction for the promise multipliers – at least in steady state.[43] But promise sequences that induce convergence in the multipliers need not be the only cases of interest.

Again, the contrast with Ramsey policy is worth emphasising. By equation (16), it is immediate that if Ramsey policy converges to a steady state with bounded multipliers, these will satisfy the relationship:

$$\lambda^m_{ss}(\sigma) + \lambda^k_{ss}(\sigma_-) = \lambda^k_{ss}(\sigma)$$

This is inconsistent with (19) whenever the multipliers are non-zero and $\beta < 1$. It follows that there must come a time period when the continuation of

[41] The latter arises because inflation means that the real effect of a unit change in $\omega_t(\sigma')$ falls.

[42] Strictly this is not a direct corollary, as cases in which either $\left[\lambda^m_{ss}(\sigma) + \lambda^k_{ss}(\sigma_-)\right] = 0$ or $\lambda^k_{ss}(\sigma) = 0$ are not covered by Proposition 9. The argument of the proof is easy to extend to these settings, however, confirming equation (19).

[43] Outcomes away from steady state are clearly crucial too. See section 9.4 for a discussion of the simple conditions that may apply when an additional symmetry requirement is imposed.

Ramsey policy has ceased to exhibit even the minimal desirability properties that characterise the set $D\left(x'_{t-1}\right)$.

8.3 Limits to dynamic variation

Proposition 9 provides insight into the long-run evolution of one particular ratio between the multipliers, for any given state σ. As discussed above, it is particularly notable that the ratio links the benefits of making a promise in t for $t+1$ with the costs of keeping an equivalent promise in t. It does not establish any necessary link between multipliers from one calendar period to the next, and the possibility of non-convergence remains open. The next two Propositions provide some limits on the extent of long-run variation in the multipliers that is consistent with time-consistently undominated policy. In both cases the main results state particular relationships that are 'close', in that deviations from them are either vanishingly small in magnitude or vanishingly infrequent as time progresses. In particular, we use the following definition:

Definition. Consider an arbitrary vector of variables $z_t \in Z$ and an arbitrary function $\phi : Z \to \mathbb{R}$. For any time period τ and any $\varepsilon > 0$, index by $n \in \{1, ..., N\}$ the set of periods t in which $\phi(z_t) \leq -\varepsilon$, with $t(n, \tau, \varepsilon)$ used to denote the time period in which the nth occurrence of this inequality arises subsequent to τ. We will say that the restriction $\phi(z_t) \geq 0$ is **almost never violated** if for all $\varepsilon > 0$ and all τ, either N is finite or:

$$\sup_n \left[t(n+1, \tau, \varepsilon) - t(n, \tau, \varepsilon)\right] = \infty$$

Notice that this definition can be applied to equality relationships of the form $\phi(z_t) = 0$ by interpreting them as two-sided inequalities.

The first Proposition we have relates to the nature of deviations from Ramsey-optimal choices.

Proposition 10. *Suppose that the policy $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ is time-consistently undominated, given some initial $x'_{s-1} \in X$. Under **difference comparability**, for Π-almost all $\sigma \in \Sigma$ the following inequality is almost never violated:*

$$\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-) \geq \lambda_{t+1}^k(\sigma) \tag{20}$$

*Under **ratio comparability**, the same applies to the inequality:*

$$\left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right]\omega_{t+1}(\sigma) \geq \lambda_{t+1}^k(\sigma)\,\omega_{t+1}(\sigma) \tag{21}$$

This result is noteworthy because of its close relationship to the Ramsey optimality condition (16). In particular, a Ramsey-optimal choice for period s satisfies (20) with equality in every period $t \geq s$. The result establishes that movements away from the Ramsey first-order condition are generally admissible in one direction, but not the other. The ratio of the cost of keeping promises, $\lambda_{t+1}^k(\sigma)$, to the *ex-ante* benefit from making them, $\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)$, can be no *greater* than it is under Ramsey policy, but it can systematically be lower.

The intuition behind this result is reasonably straightforward, and clarifies the trade-offs that time-consistently undominated policy allows. When inequality (20) is violated, there is no disagreement between policymakers in different time periods about the appropriate marginal direction in which to move $\omega_{t+1}(\sigma)$. *Ex-post*, it is always desirable to relax promise-keeping constraints, and so the policymaker in $t+1$ would prefer a reduction in $\omega_{t+1}(\sigma)$ regardless of the inequality. When the inequality is violated, this is an opinion that is shared by policymakers prior to $t+1$, and so there is a profitable improvement from the perspective of all time periods. Thus, by Proposition 4, the original promise sequence is dominated.

The next Proposition considers the relative size of the multiplier values across states.

Proposition 11. *Suppose that the policy* $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ *is time-consistently undominated, given some initial* $x'_{s-1} \in X$*. There exists a sequence of scalar values* $\{\alpha_t\}_{t\geq s}$ *with* $\alpha_t \in [0,1]$ *for all* t*, such that for* Π*-almost all* $\sigma \in \Sigma$*, under* ***difference comparability*** *the following equality is almost never violated:*

$$\alpha_t \left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right] = \lambda_{t+1}^k(\sigma) \tag{22}$$

and under ***ratio comparability*** *the following equality is almost never violated:*

$$\alpha_t \left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right] \omega_{t+1}(\sigma) = \lambda_{t+1}^k(\sigma)\, \omega_{t+1}(\sigma) \tag{23}$$

This result is interesting because of the substantial restriction that condition (22) places on cross-sectional variation in the promise multipliers. Proposition 10 established that for any given σ, the ratio $\frac{\lambda_{t+1}^k(\sigma)}{\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)}$ could not systematically exceed one. This ratio captures the shadow cost of keeping a promise in state σ at time $t+1$, versus the *ex-ante* shadow benefit from making it. A value of one is the extreme case satisfied by Ramsey policy. Proposition 11 establishes that whatever value this ratio takes for any particular $\sigma \in \Sigma$, it should take the *same* value for all other states. That is, the value of α_t is independent of σ. Thus for any given period, there is commonality across states in the extent of the departure from Ramsey policy.

This makes intuitive sense. A time-consistently undominated policy differs from Ramsey in its treatment of the relative preferences of policymakers at different points in time. There is no reason for variation in this difference across states. Whenever this variation exists, there are gains to all from smoothing it out. Indeed, the state-invariant constant α_t can be interpreted as a form of intertemporal Pareto weight, measuring the relative weight given to preferences prior to $t+1$ versus preferences in $t+1$ when determining the size of promises for $t+1$. Proposition 9 implies this value will generally live in the neighbourhood of the discount factor β when policy is chosen to be time-consistently undominated.

The final Proposition in this section is useful for contrasting time-consistently undominated policy with the outcome of period-by-period discretionary choice. We have the following.

Proposition 12. *Suppose that the policy* $(\bar{\mathbf{x}}_{\mathbf{s}}', \bar{\mathbf{a}}_{\mathbf{s}}')$ *is time-consistently undominated, given some initial* $x_{s-1}' \in X$, *and that there is a positive measure of* $\sigma \in \Sigma$ *for which under* ***difference comparability*** *the following inequality is almost never violated:*

$$\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-) \geq \varepsilon$$

for some $\varepsilon > 0$. *Then for* Π*-almost all* σ *in this set, the condition* $\lambda_t^k(\sigma) = 0$ *can only occur with vanishingly low frequency. Specifically, there does not exist a* $T < \infty$ *such that for every* $t \geq s$ *there is a* $t + \tau$ *with* $\lambda_{t+\tau}^k(\sigma) = 0$ *and* $\tau < T$.

Under ***ratio comparability*** *the same applies, with the inequality amended to:*

$$\left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right]\omega_{t+1}(\sigma) \geq \varepsilon$$

In words, this says that if the marginal benefits of making promises are bounded above zero almost all of the time, then the marginal costs of keeping them should be likewise. If this were not true, then at every point in time there would exist beneficial future promises that are not being made, even though they impose zero marginal costs.

This result is of interest for two reasons. First, it demonstrates that a 'discretionary' approach to policy design is dominated under our criteria except in trivial cases. Discretionary policy design, also commonly referred to as Markovian or time-consistent policy choice, assumes that there is no commitment technology that can be used by past policymakers to bind their successors. This extends to ruling out history dependence in all agents' choice strategies, so that trigger strategies cannot be used to support even modest commitments. A possible implication of this equilibrium concept is that the resulting policy need not solve Problem 1 for the promises that it induces, since there are strategic incentives to manipulate the choice of state variables through time.[44] More significantly for present purposes, if past promises can never bind, $\lambda_t^k(\sigma) = 0$ must always hold for all t and all σ. The Proposition highlights that there are always time-consistent improvements relative to this, except in the trivial case that there are zero benefits to making promises.

The second implication of Proposition 12 relates back to Proposition 4, 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, at least so long as promises matter.

8.4 Sufficiency

Propositions 9 to 12 place necessary restrictions on the long-run evolution of the promise multipliers when policy is time-consistently undominated. We have

[44] See, for instance, Klein, Krusell and Ríos-Rull (2008).

discussed the important classes of policy that these conditions rule out, but for more constructive purposes we need sufficiency results. This section provides general conditions that guarantee that a policy never comes to be dominated. Again, the focus initially is on the most general statements possible, restricting to more applicable expressions subsequently.

Proposition 13. *Consider a policy $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ that solves Problem 1 for the promise sequence that it induces, $\bar{\omega}'_{\mathbf{s}}$. The continuation of this policy $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ will belong to $D\left(x'_{t-1}\right)$ for all $t \geq s$ provided the following are true:*

1. *The value function $V\left(\bar{\omega}_{\mathbf{s}}; x_{s-1}\right)$ is concave in $\bar{\omega}_{\mathbf{s}}$.*
2. (a) *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:*

$$\lim_{t\to\infty}\left[\frac{\lambda_{t+1}^k(\sigma)}{\alpha_t\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}\right]=1 \tag{24}$$

where the rate of convergence is at least linear.

(b) *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**:*

$$\underline{K} \leq \prod_{t=\tau}^{r-1}\frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]} \leq \bar{K} \tag{25}$$

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

$$\underline{K} \leq \prod_{t=\tau}^{r-1}\frac{\lambda_t^k(\sigma)\,\omega_t(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]\omega_{t+1}(\sigma)} \leq \bar{K} \tag{26}$$

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.

The second part of the Proposition provides restrictions on the multipliers that closely mirror the necessary restrictions provided in Propositions 9 to 11.[45] A policy is time-consistently undominated provided its multipliers from one period to the next converge to satisfying a common ratio α_t across states, and provided the compounded ratio of promise multipliers in (25) remains bounded. Both of these conditions were discussed in detail above. The first can be read as

[45] The main differences are that the equality in (a) is imposed in every period from τ onwards for all $\sigma \in \Sigma$, rather than the 'almost everywhere' specification in Propositions 10 and 11, and that the bounds in part (b) are slightly tighter than those applied in Proposition 9. In both cases this is done to simplify the statement of the sufficiency requirements. Thus it would be possible to extend Proposition 13 to narrow the gap further, but at a notational cost.

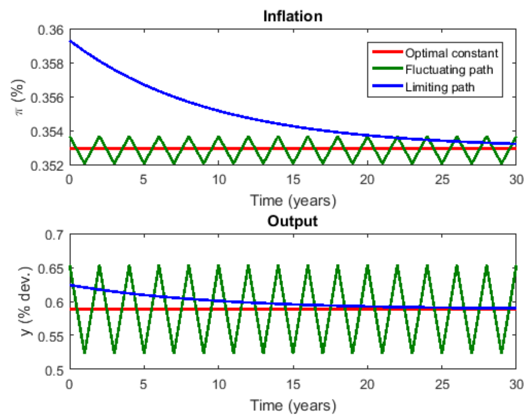

Figure 12: Inflation bias: Three time-consistently undominated policies

a generalisation of the Ramsey dynamic multiplier condition (16), allowing for a lower willingness to bear the costs of keeping promises, relative to the gains from making them. Note again that the ratio α_t need not itself converge to any fixed value. The second condition loosely implies that the relative costs from keeping promises are balanced with the benefits from making them as time progresses. Viewed in period t, $\lambda_t^k(\sigma)$ is the marginal cost of increasing $\omega_t(\sigma')$ by a unit, and $\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]$ is the marginal benefit from increasing $\omega_{t+1}(\sigma')$ by a unit.

8.4.1 Multiple possibilities

A consequence of the freedom in α_t permitted under Proposition 13 is that it will commonly 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. Figures 8 and 9, reproduced below, highlighted that although just one *constant* path was time-consistently undominated, there are many time-varying paths with this property. For the examples charted, the fact that the policies are time-consistently undominated is verified by confirming that conditions 2(a) and 2(b) of Proposition 13 are satisfied numerically.[46]

The possibility of limit cycles in the policy choice is of analytical interest, but arguably of little practical relevance. More problematic in making an active policy selection is the fact that time-consistently undominated policies are not restricted in their transition to steady state. As the example illustrates, any two promise sequences converging to the same, time-consistently undominated steady-state outcome will be admissible.

When the problem as a whole is stationary, as here, the constant outcome still appears a 'natural' selection to make among the set of convergent choices.

[46] Concavity of the value function for this example is easy to confirm.

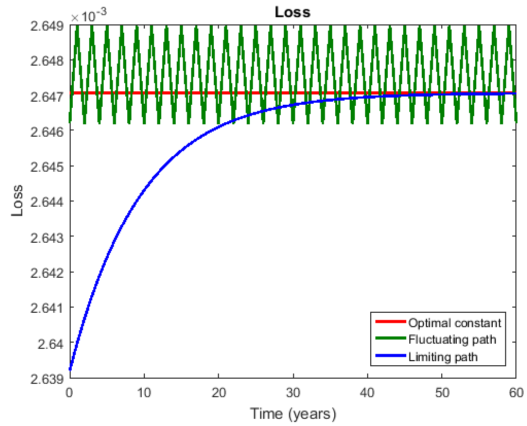

Figure 13: Inflation bias: Loss for three time-consistently undominated policies

This becomes far less immediate, however, when there are non-trivial dynamics in the underlying economic system. If, for instance, the capital stock is evolving over time, it seems at least possible that the chosen promise sequence ought to be evolving with it. Assuming that a steady state is ultimately reached, this means that promises must follow some sort of transition dynamic. But then how to select a unique transition? Can promises still be chosen in a way that is meaningfully 'symmetic' through time, once a constant choice has been surrendered? The next section provides the tools to address this question.

9 An equivalent representation

The focus of attention so far has been on the problem of choosing a policy that is time-consistently undominated, given the normative criteria set out. In this section we show that a policy that solves this problem will have a parallel interpretation that is central in allowing our approach to be operationalised, and to address the multiplicity problem just outlined. Specifically, a time-consistently *undominated* choice of promises is also a time-consistently *optimal* choice of promises in a problem with suitably restricted degrees of freedom. The relevant cross-restriction applies over time, so that any given promise choice for t will only be possible in conjunction with specific choices in periods other than t.

This is useful for two purposes. First, the fact that a policy must be time-consistently optimal in the restricted-dimensional problem provides a clear route to obtaining simple first-order optimality conditions that can be used to characterise and compute time-consistently undominated policy. Much of the foregoing discussion has related to the set of undominated sequences D – a large, abstract set that is near-impossible to visualise or compute. The alternative specification of the problem provides a much clearer route to operationalising and computing policies that belong to D indefinitely.

Second, the time-consistent choice problem makes it possible to provide a general definition of policies that are 'symmetric' over time. As explained above, there may be many dynamic allocations that are time-consistently undominated, and the existing normative conditions give no guidance on how to choose among them. Yet we show that there will generally exist a unique allocation that can be recovered as the time-consistent outcome of a problem that is symmetric in the way that promise choice is restricted through time. Symmetry is defined in a way that is invariant to equivalent representations of the promises. The symmetric undominated policies that follow from this problem generally have a simple and intuitive character that is appealing for communication purposes.

The idea can be easiest to see by reference to a specific example. In the inflation bias example of Section 2, consider the restricted problem of choosing a *constant* rate of inflation to implement from period s onwards. Formally this problem has one degree of freedom, with a cross restriction limiting the inflation rate in period t to be equal to the inflation rate in period s for all $t > s$. Now, it is trivially true that the continuation of this constrained-optimal policy from period $t > s$ onwards coincides with the solution to the problem of finding the optimal constant rate of inflation from t onwards. Thus even though the problem of finding an unrestricted sequence of inflation rates from s onwards is time inconsistent, the restricted problem of choosing a constant inflation rate induces agreement through time.

The point of this section is to generalise this insight. Policies that are time-consistently undominated quite generally have a representation as the time-consistently optimal choice from some restricted dimensional choice problem. In addition, as in the 'constant inflation rate' example, some of these policies may be notable for their symmetry over time.

9.1 Preliminaries

9.1.1 Restricted-dimensional choice

We will consider the problem of choosing a promise sequence $\bar{\omega}_{\mathbf{s}}$ from some restricted-dimensional subspace of $\mathcal{W}$, where this subspace is defined parametrically by reference to a benchmark sequence $\bar{\omega}'_{\mathbf{s}} \in \mathring{\Omega}(x_{s-1})$ and a set of possible vector movements away from $\bar{\omega}'_{\mathbf{s}}$. This subsection provides the basic mathematical apparatus for doing this. 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 permissable rescalings of the promise values. Once more, this requires the cases of difference comparability and ratio comparability to be treated separately.

We focus first on the case of difference comparability. For all $\sigma \in \Sigma$, fix a sequence of strictly positive scalars $\{\delta_t(\sigma)\}_{t \geq s}$, with $\delta_t(\sigma) \in [\underline{\delta}, \bar{\delta}]$ for all t and σ, given some $\underline{\delta}$ and $\bar{\delta}$ satisfying $0 < \underline{\delta} \leq \bar{\delta} < \infty$. Given the benchmark sequence $\{\omega'_t(\sigma)\}_{t \geq s}$ and these 'slope' values $\{\delta_t(\sigma)\}_{t \geq s}$, we will focus on the one-dimensional set of of σ-contingent promise sequences $\{\omega_t(\sigma)\}_{t \geq s}$ that can be traced out by varying the common scalar $\theta(\sigma)$ in the equation:

$$\omega_t(\sigma) = \omega'_t(\sigma) + \theta(\sigma)\,\delta_t(\sigma)$$

for all $t \geq s$. The ability to select a value for $\theta(\sigma)$ implies the ability to move the promises $\{\omega_t(\sigma)\}_{t\geq s}$ jointly along exactly one dimension, with the marginal effect of a unit change to $\theta(\sigma)$ on $\omega_t(\sigma)$ given by $\delta_t(\sigma)$. Thus it is possible to increase $\omega_s(\sigma)$ by one unit only if $\omega_t(\sigma)$ is simultaneously increased by $\frac{\delta_t(\sigma)}{\delta_s(\sigma)}$ units for all $t > s$. Under difference comparability this relative change in the promise values for different time periods is well defined, independently of normalisation to the levels of $\omega_s(\sigma)$ and $\omega_t(\sigma)$.

The corresponding construction under ratio comparability is a benchmark sequence of promises $\{\omega'_t(\sigma)\}_{t\geq s}$ and a specified vector of joint proportional movements away from $\{\omega'_t(\sigma)\}_{t\geq s}$, again characterising a unidimensional set of promise sequences $\{\omega_t(\sigma)\}_{t\geq s}$. For this, let $\{\delta_t(\sigma)\}_{t\geq s}$ again be a sequence of positive scalars satisfying uniform upper and lower bounds, and $\theta(\sigma)$ a time-invariant parameter[47] such that the set of permitted $\{\omega_t(\sigma)\}_{t\geq s}$ for all $t \geq s$ $\omega_t(\sigma)$ satisfies:

$$\omega_t(\sigma) = \omega'_t(\sigma)\exp\{\theta(\sigma)\,\delta_t(\sigma)\}$$

Thus the effect of increasing $\theta(\sigma)$ by one unit at the margin will be to increase $\omega_t(\sigma)$ by $\delta_t(\sigma)\,\omega_t(\sigma)$ units.

9.1.2 Notation

Irrespective of the form of comparability, we will define $\bar{\delta}_{\mathbf{s}}$as the array of $\delta_t(\sigma)$ sequences across date-states:

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

and θ will denote an array of $\theta(\sigma)$ choices across σ:

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

Notation in this section is simplified by defining standard operations on the objects θ, $\bar{\delta}_{\mathbf{s}}$ and $\bar{\omega}_{\mathbf{s}}$ elementwise. Thus we will write $\theta\bar{\delta}_{\mathbf{s}}$ to denote the array obtained by elementwise multiplication:

$$\left\{\{\theta(\sigma)\,\delta_t(\sigma)\}_{\sigma\in\Sigma}\right\}_{t\geq s}$$

$\exp\left\{\theta\bar{\delta}_{\mathbf{s}}\right\}$ to denote the array:

$$\left\{\{\exp\{\theta(\sigma)\,\delta_t(\sigma)\}\}_{\sigma\in\Sigma}\right\}_{t\geq s}$$

and similarly.

[47] More generally $\theta(\sigma)$ can belong to a bounded subset of the complex numbers $\mathbb{C}$. This allows for $\omega_t(\sigma)$ to take a different sign to $\omega'_t(\sigma)$.

Finally, it is useful to be able to refer in general to the choice of $\bar{\omega}_{\mathbf{s}}$ as a function of θ, given $\bar{\delta}_{\mathbf{s}}$ and $\bar{\omega}'_{\mathbf{s}}$, irrespective of the form of comparability assumed. Thus, we let:

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

9.2 A restricted-dimension problem

We will consider the following problem:

Problem 2. (Restricted Promise Choice)

$$\max_{\theta} V\left(\bar{\omega}_{\mathbf{s}}\left(\theta;\bar{\omega}'_{\mathbf{s}},\bar{\delta}_{\mathbf{s}}\right);x_{s-1}\right)$$

given $\bar{\omega}'_{\mathbf{s}}$ and $\bar{\delta}_{\mathbf{s}}$.

Assuming that it exists, the value of θ that solves this problem is denoted θ^*, with the resulting promise vector $\bar{\omega}^*_{\mathbf{s}}:=\bar{\omega}_{\mathbf{s}}\left(\theta^*;\bar{\omega}'_{\mathbf{s}},\bar{\delta}_{\mathbf{s}}\right)$, 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 $\mathring{\Omega}\left(x_{s-1}\right)$, the interior of $\Omega\left(x_{s-1}\right)$. Then by standard calculus a necessary optimality condition for the choice of $\theta^*\left(\sigma\right)$ is:[48]

$$\sum_{r=s}^{\infty}\beta^{r-s}\left\{\lambda^k_r\left(\sigma\right)\delta_r\left(\sigma\right)-\beta\left[\lambda^m_r\left(\sigma\right)+\lambda^k_r\left(\sigma_-\right)\right]\delta_{r+1}\left(\sigma\right)\right\}=0 \tag{27}$$

for the case of difference comparability and:

$$\sum_{r=s}^{\infty}\beta^{r-s}\left\{\lambda^k_r\left(\sigma\right)\omega_r\left(\sigma\right)\delta_r\left(\sigma\right)-\beta\left[\lambda^m_r\left(\sigma\right)+\lambda^k_r\left(\sigma_-\right)\right]\omega_{r+1}\left(\sigma\right)\delta_{r+1}\left(\sigma\right)\right\}=0 \tag{28}$$

for the case of ratio comparability. If the value function is concave in $\bar{\omega}_{\mathbf{s}}$ then this condition holding for all $\sigma\in\Sigma$ is also sufficient for Problem 2 to be solved.

Heuristically the main interest in Problem 2 arises from the possibility that the optimal θ^* may stay constant through time, given the chosen $\bar{\delta}_{\mathbf{s}}$ and $\bar{\omega}'_{\mathbf{s}}$ sequences. If this were the case, (27) or (28) would hold for all possible initial periods, not just the original s. Applying this in $s+1$ implies that the forward sum will cancel, leaving a single within-period restriction that must hold for period s, and identically for all $t\geq s$:

$$\lambda^k_t\left(\sigma\right)\delta_t\left(\sigma\right)-\beta\left[\lambda^m_t\left(\sigma\right)+\lambda^k_t\left(\sigma_-\right)\right]\delta_{t+1}\left(\sigma\right)=0 \tag{29}$$

under difference comparability or:

$$\lambda^k_t\left(\sigma\right)\omega_t\left(\sigma\right)\delta_t\left(\sigma\right)-\beta\left[\lambda^m_t\left(\sigma\right)+\lambda^k_t\left(\sigma_-\right)\right]\omega_{t+1}\left(\sigma\right)\delta_{t+1}\left(\sigma\right)=0 \tag{30}$$

[48] Note that the boundedness of $\delta_r\left(\sigma\right)$ ensures that this sum is convergent.

with ratio comparability. The stationarity of the problem then implies that equivalent conditions will additionally hold for all periods $t \geq s$.

Suppose that $\bar{\omega}_{\mathbf{s}}$ indeed satisfies the relevant version of this restriction for all t. Then consider the following product ratio under difference comparability:

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

The boundedness requirements 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 $\bar{\omega}_{\mathbf{s}}^*$ will satisfy sufficiency condition 2(b) in Proposition 13, at least for this value of σ. This is at least suggestive of a link between the time-consistent choice of θ and time-consistent undominance. Under ratio comparability a similar argument can be applied, this time for the ratio:

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

Again, boundedness in the $\delta_r(\sigma)$ coefficient implies boundedness in the ratio on the left-hand side, which is precisely the condition required for time-consistently undominated policy.

More formally, the equivalence between time-consistent solutions to Problem 2 and time-consistently undominated policies is summarised in the following Proposition.

Proposition 14. *Suppose the value function $V(\bar{\omega}_{\mathbf{s}}; x_{s-1})$ is concave in $\bar{\omega}_{\mathbf{s}}$. Then an allocation $(\bar{\mathbf{x}}_{\mathbf{s}}*, \bar{\mathbf{a}}_{\mathbf{s}}*)$, inducing promises $\bar{\omega}_{\mathbf{s}}*$, satisfies sufficiency condition 2(b) of Proposition 13 if and only if there is a sequence $\bar{\delta}_{\mathbf{s}}$ such that $\bar{\omega}_{\mathbf{s}}*$ solves Problem 2 recursively, given $\bar{\delta}_{\mathbf{s}}$ and $\bar{\omega}_{\mathbf{s}}' = \bar{\omega}_{\mathbf{s}}*$.*

This follows from the foregoing discussion, and a formal proof is omitted.[49]

The Proposition does not directly imply equivalence between time-consistently undominated policies and recursive solutions to Problem 2, since recursive solutions to Problem 2 need not generically satisfy sufficiency condition 2(a) of Proposition 13 – only 2(b). To verify whether condition 2(a) is additionally satisfied requires checking the properties of any candidate solution, and since the condition places a convergence requirement on a multiplier ratio across states, this appears quite onerous. For practical purposes, however, the following corollary will have wide applicability:

Corollary 3. *Suppose the value function $V(\bar{\omega}_{\mathbf{s}}; x_{s-1})$ is concave in $\bar{\omega}_{\mathbf{s}}$, and that difference comparability applies. If an allocation $(\bar{\mathbf{x}}_{\mathbf{s}}*, \bar{\mathbf{a}}_{\mathbf{s}}*)$, inducing promises $\bar{\omega}_{\mathbf{s}}*$, solves Problem 2 recursively, given some $\bar{\delta}_{\mathbf{s}}$ and $\bar{\omega}_{\mathbf{s}}' = \bar{\omega}_{\mathbf{s}}*$, and induces convergence in the intertemporal multipliers to steady-state values $\lambda_{ss}^m(\sigma) > 0$*

[49] Concavity of the value function is needed to ensure that global optimality follows from local first-order conditions.

and $\lambda_{ss}^{k}(\sigma) > 0$ *for* Π*-almost all* $\sigma \in \Sigma$*, then* $(\bar{\mathbf{x}}_{\mathbf{s}}*, \bar{\mathbf{a}}_{\mathbf{s}}*)$ *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* $\sigma \in \Sigma$*.*

As discussed above, steady-state convergence is certainly not a necessary property for a time-consistently undominated policy, but it is an extremely simple property to verify when it does arise. Indeed, the most straightforward computational approach to solving for a time-consistently undominated policy will often be first to solve for a steady-state allocation, and then to assess convergence dynamics towards this steady state. This imposes the convergence property directly through the solution technique, and so full sufficiency will be guaranteed provided the allocation recursively solves Problem 2, and the value function is quasi-concave in the promises. Both of these features are easily verified.

When promises are ratio comparable the condition additionally requires convergence in the promise values. By itself this is always possible to ensure by normalisation, under any allocation that does not imply convergence to the zero value. But a different normalisation for the promises implies a different set of values for the multiplier terms over time. This is because the multipliers capture the slope of V with respect to a *unit* change in $\omega_t(\sigma)$, and ratio permits arbitrary rescaling of these units over time. Thus the requirement in this case is that the multipliers converge *for any normalisation that implies convergence in the promises.*

9.3 Interpretation

The duality between a time-consistently undominated promise sequence and a time-consistent choice of θ has parallels in more conventional economic domains. Consider a market equilibrium in a Walrasian endowment economy that satisfies the first welfare theorem. By definition, this equilibrium is efficient. Given the price vector that supports the equilibrium, moving allocations away from an initial endowment towards the equilibrium position is beneficial to all agents. Given the restricted-dimensional choice set implied by their budget constraints, all agents are content to trade up to the equilibrium point, and no further.

Here, the chosen value for $\bar{\omega}_{\mathbf{s}}^*$ is analogous to an efficient allocation. It is a promise sequence with the property that no other sequence is strictly preferred at every point in time. Allowing choice in each time period within the restricted-dimensional space characterised by θ is the analogue of allowing each agent in a Walrasian problem to trade away from their initial endowment according to a given price ratio. The finding that $\bar{\omega}_{\mathbf{s}}^*$ solves Problem 2 recursively at some θ^*, when $\bar{\delta}_{\mathbf{s}}$ and $\bar{\omega}_{\mathbf{s}}'$ are appropriately specified, is the equivalent of finding that a Walrasian equilibrium has a supporting price vector, and a set of endowments that are consistent with trade to it.

The analogy can further shed light on condition 2(a) in Proposition 13, which is an additional requirement to confirm that a policy is time-consistently un-

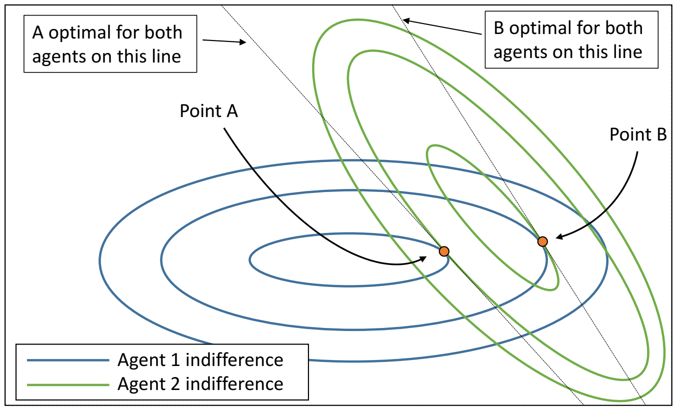

Figure 14: Agreement along a restricted dimension need not imply Pareto efficiency

dominated. Figure 14 illustrates indifference curves for two distinct agents in an arbitrary two-dimensional space – akin to a classic Edgeworth box. Preferences are assumed to be non-monotonic for each of the two agents, with preferred options contained within the elliptical indifference curves. Consider point A first. This is a Pareto efficient choice with respect to the two agents' preferences, and the straight line through the point represents a supporting price vector. Despite the general disagreement in their preferences, if either agent were given the choice of finding an optimal allocation along this line, they would choose point A. This restricted problem is a two-dimensional analogue to Problem 2, which is likewise solved by the same choice of allocation for the policymaker at each distinct point in time – even though preferences over promises are generally time-varying.

Point B is a different case. Like point A, it is an optimal allocation for both agents along the line that is drawn through it: indifference curves for both agents are tangential to this line at B. However, unlike point A this is clearly not a Pareto efficient allocation. Both agents would gain by moving to the 'south-west' of the point. More generally, agreement along one dimension is implied by any Pareto efficient allocation, but the converse need not hold. An allocation can be best for all agents in a restricted-dimensional subspace, but not be Pareto efficient.

Condition 2(a) in Proposition 13 is additionally needed for this reason. It ensures that an allocation that is optimal in every period along a single dimension also has the property that no other allocation would be strictly preferred from the perspective of every point in time. This amounts to ensuring that – at least at the limit – there is a single relative Pareto weight α_t placed on the preferences of policymakers in time period t and earlier, when assessing promises that are to be kept in period $t+1$ from the perspective of some earlier time $s \leq t$. The value of α_t is common across all stochastic draws σ, and must lie in

$(0, 1)$. The weight on the preferences of the policymaker in $t + 1$ is $(1 - \alpha_t)$. If there was no convergence to a weight structure of this kind, there would exist changes to the dynamic allocation with the potential to make the policymaker strictly better off at every point in time.

9.4 Symmetry

A useful feature of Problem 2 is that it provides a route to selecting a unique symmetric policy over time. The appeal of symmetry is intrinsically connected to the basic motivation for seeking time-consistently undominated policy. That is, given a policy problem exhibiting time inconsistency, is it possible to select an allocation in a way that is both normatively appealing, and invariant to the time period in which selection takes place? As the inflation bias example of Section 2 illustrated, it is quite possible to find policies that are time-consistently undominated, but exhibit fluctuations and cycles at the limit that are unrelated to any change in the underlying economic environment. The idea behind a symmetry condition is to rule out these sorts of dynamics, identifying time-consistently undominated allocations that operate 'in the same manner' in all time periods.

There are two main difficulties in doing this. The first is that most examples of interest will not feature a stationary economic environment through time. In models with capital accumulation dynamics, for instance, accumulation or depletion of the capital stock over time will change the basic resource constraints that feature in the problem, and this means it will not be feasible for every aspect of the allocation to be invariant through time. Thus the question arises: 'symmetry along what dimension'? This is where well-defined comparability properties in the constraint functions will become essential, since it would not make sense to impose dynamic symmetry along a dimension that was sensitive to renormalisation.

The second difficulty relates to an option that is not easily available. If time inconsistency problems are problems that arise because policymakers at different points in time have different preferences over promises, it might seem natural to resolve this in period s by placing different Pareto weights on all of the policymakers' value functions from s onwards, and finding a promise choice that maximises the resulting weighted sum. A 'symmetric' choice might then be one that puts equal, or perhaps geometrically decaying, weights on every period. The difficulty in applying this approach to the present problem is that it does not resolve time inconsistency. The benefits from making promises always arise in periods prior to those promises being kept, so that even the Pareto-weighted choice for s onwards would never allow promises in s to restrict choice meaningfully.[50] But then re-applying the same technique in any $t > s$ would imply no promises in t could be binding. So either the method could not be applied recursively, or it would generate a policy in violation of the necessary condition stated in Proposition 12.

[50]That is, $\lambda_s^k(\sigma)$ would be zero for all σ.

Problem 2 gives the scope to define a set of vector changes to the promises over time that are symmetric, where symmetry in turn is defined with respect to a vector movement that has economic meaning in the given problem. In terms of the single instrument θ, this means finding a promise sequence that is time-consistently optimal along a dimension that provides equal scope to change outcomes in each period. Given the re-scaling possibilities for the promise values themselves, this is as close as it is possible to get to treating each period's policymaker equally when selecting a promise vector. Alternative approaches to symmetry could easily be imagined – for instance, choosing a constant value for $\omega_t(\sigma)$ across t – but these would deliver policies that varied in equivalent representations of the constraint functions.

Irrespective of its precise definition, whether symmetry should be viewed as a desirable normative property *per se* is debatable, but at the very least it will provide a useful benchmark among the multiple possibilities for time-consistently undominated choice. Moreover, since the purpose of the solution approach in this paper is to arrive at policy rules that follow from repeated application of the same normative techniques, symmetric treatment of time periods is a necessary step.

In practice the restriction means that the $\delta_t(\sigma)$ terms will be required to deliver either common level changes in $\omega_t(\sigma)$ for all t as $\theta(\sigma)$ varies, or common proportional changes. Consider first a case in which the level changes in the promises are defined relative to one another. 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 problem in period s is:

$$\sum_{r=s}^{\infty} \beta^{r-s} \left\{ \lambda_r^k(\sigma) - \beta \left[\lambda_r^m(\sigma) + \lambda_r^k(\sigma_-) \right] \right\} = 0 \tag{31}$$

and once more, if the choice of $\bar{\omega}_{\mathbf{s}}$ is able to solve the same problem recursively in every period, this delivers the following restriction for all t and all σ:

$$\lambda_t^k(\sigma) - \beta \left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-) \right] = 0 \tag{32}$$

Equivalently, if proportional changes in the promises are defined then a symmetric version of Problem 2 would require that $\delta_s(\sigma)\,\omega_s(\sigma) = \delta_t(\sigma)\,\omega_t(\sigma)$ for all $t > s$ and all $\sigma \in \Sigma$, so 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:

$$\lambda_t^k(\sigma)\,\omega_t(\sigma) - \beta \left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-) \right] \omega_{t+1}(\sigma) = 0 \tag{33}$$

9.5 Sufficiency again

A sequence of promises that implements the relevant version of condition (32) or (33) in every period will still not necessarily imply time-consistently undominated policy, since the convergence requirement 2(a) in Proposition 13 may not have been met.[51] A useful feature of a symmetric policy, however, is that it

[51] The value function may also fail to be quasiconcave in the promises, but this causes general problems for sufficiency, irrespective of the solution concept.

will be consistent with steady state being achieved. If this steady state extends to multiplier values, it will be much more straightforward to confirm that the required convergence has occurred. Consider equation (32), for instance. If convergence in the multiplier values occurs, we will have:

$$\lim_{t\to\infty}\left[\lambda_t^k(\sigma)\right]=\lambda_{ss}^k(\sigma)=\beta\lim_{t\to\infty}\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]$$

Thus the object in condition (24) will indeed converge, with $\alpha_t=\beta$. An equivalent result would obtain were the symmetry requirement imposed on proportional promise changes.

Without symmetry, this sufficiency result would be far more difficult to check. The requirement is that there should be a common α_t across states σ in the converging object:

$$\frac{\lambda_{t+1}^k(\sigma)}{\alpha_t\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}$$

Checking this across all σ in the absence of convergence is a far greater challenge.

In short, checking 2(a) in the sufficiency requirements will be near-trivial in most applications of the symmetric solution – it needs only that convergence in the multipliers should occur. This gives a more practical reason for avoiding asymmetric solutions.

10 Applications

[To be added.]

11 Relation to literature

11.1 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, 2016) has provided alternative devices for representing the Ramsey problem in recursive form.[52] Our use of promise values in characterising Problem 1 above has similarities with the method of Abreu et al., though our aim is not to provide a recursive representation, and we define promises in sequence space. Our

[52] Though Abreu et al. were writing 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).

characterisation results, by contrast, are stated in terms of the promise multipliers whose use Marcet and Marimon popularised, and are easiest to interpret by comparison with their characterisation of Ramsey choice.

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, sometimes referred to as 'time-consistent' policy.[53] These allow no scope for promises to bind past choice, though strategic incentives to influence future decisions can affect the choice of endogenous state variables. Consistent with the original results of Kydland and Prescott, welfare outcomes are generally undesirable, with alternative feasible strategies delivering improvements in every time period. In terms of the dominance ordering that we use, these outcomes are preference-dominated.

A smaller, though highly influential, literature focuses on history-contingent reputational equilibria.[54] This 'sustainable plans' approach characterises the set of policies that can be supported by appropriate trigger strategies in an infinite horizon. Unlike Markov equilibria, policymakers have the capacity to make and keep some promises. This is because the threat of reversion to an inferior equilibrium can outweigh the *ex-post* cost of promise-keeping. A common feature of this literature is indeterminacy in the equilibrium outcome. There will generally be many supportable policy strategies that improve on a Markov outcomes, any one of which could be specified. The time-consistently undominated policy that we emphasise will often belong to this set,[55] and – if so – our work could be interpreted as a normative case for a particular focal point in the equilibrium set.

11.2 The timeless perspective

As discussed in the introduction, the basic problem of finding a time-consistent normative policy choice in Kydland and Prescott problems has been most directly studied 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.[56] It is often used to obtain time-invariant 'target criteria' that characterise a trade-off between inflation and measures of real activity. The approach is usually de-

[53] Examples 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) and Bianchi and Mendoza (2013).

[54] Chari and Kehoe (1990) and Atkeson (1991) were pioneering early papers. See also the specific discussion of the dynamic social insurance literature below.

[55] See the discussion of Sleet and Yeltekin (2006) below.

[56] Recent 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).

fended as a way to overcome the 'initial-period problem' – i.e., to ensure that policy choices in early time periods do not take a different form from subsequent choices.

Our results sound a note of caution about the timeless perspective. As we have shown, the long-run continuation of Ramsey policy can generically be dominated by an alternative selection that delivers strictly higher welfare at every point in time. This makes it a poor basis for time-consistent normative choice.[57] Intuitively, a policy that is designed to maximise welfare for the first-period's policymaker will generally allow welfare in later periods to be reduced if sufficient initial benefits accrue as a result. A timeless strategy accepts these costs without the benefits.

More positively, our approach may help to address some of the difficulties that are encountered when implementing timeless perspective policy. It is widely acknowledged in the literature that timeless policy is harder to study analytically in stochastic environments where the non-stochastic steady state is not efficient.[58] Specifically, the common strategy of finding an approximate linear-quadratic representation of the optimal policy problem encounters problems when there are first-order incentives to depart from from the steady state. Our method will recommend a different steady state from Ramsey policy, and it will be a steady state that the policymaker will not wish to depart from in the absence of shocks – provided the policymaker seeks a time-consistently undominated strategy. This seems likely to allow more straightforward characterisations of policy trade-offs, even in inefficient economies.

11.3 Dynamic social insurance

11.3.1 Normative literature

Our paper also contributes to the dynamic social insurance literature. Social insurance can be studied in the context of a dynamic asymmetric information problem, either of the moral hazard (hidden action) variety, or the screening (hidden information) variety.[59] A surprising, and unsettling, feature of many

[57]The desirability of the timeless approach has been questioned in the context of a linear-quadratic New Keynesian problem by Blake (2001), who shows that an alternative strategy delivers higher expected welfare in steady state. This alternative exactly coincides with our time-consistently undominated policy for Blake's example. Damjanovic, Damjanovic and Nolan (2008) refer to Blake's method as 'unconditional' welfare optimisation, and provide a general algorithm for maximising steady-state welfare. The time-consistently undominated policy that we present does not maximise steady-state welfare in the presence of state variables, and so differs from this.

[58]This is the focus of Benigno and Woodford (2005). In economies with heterogeneous agents, 'efficient' means 'first-best' under the chosen welfare criterion. Benigno, Eggertsson and Romei (2016) and Andrés, Arce and Thomas (2013) are recent examples of papers where this issue was discussed.

[59]Green (1987) and Thomas and Worrall (1990) provided pioneering studies of the dynamic moral hazard problem. Atkeson and Lucas (1992) introduced a dynamic model with hidden information about consumption preferences. Work by Golosov, Kocherlakota and Tsyvinski (2003) and Werning (2002) extended the classic Mirrlees hidden ability model to an infinite-horizon setting, generalised the earlier work of Diamond and Mirrlees (1978).

such problems is that a Ramsey-optimal plan involves 'immiseration' at the limit as time progresses.[60] That is, the consumption of almost all agents optimally tends to some lower bound, generally zero. This is true under standard assumptions about social preferences, the utility function and the market real interest rate.[61] It occurs because of an extreme growth in inequality rather than resource exhaustion.[62] Inequality increases through time because agents must be provided with incentives to produce, or reveal their type, in each successive period. *Ex-post* consumption inequality achieves this, and is allowed to accumulate through time – consistent with the martingale property of Ramsey multipliers, discussed in detail in Section 8.1.

Work by Phelan (2006) and Farhi and Werning (2007) recognises that immiseration is a troubling conclusion *per se*, and investigates options for overcoming it.[63] Like our paper, the approach of these authors is explicitly normative, with the assumption of a perfect commitment device. In both cases, the essential strategy proposed is to raise the societal discount factor – to unity in the case of Phelan (2006), and to some number in $(\beta, 1)$ in the case of Farhi and Werning (2007). This is justified on first principles as identifying an alternative position on the intergenerational Pareto frontier. Phelan (2006) goes so far as to invoke Rawls (1971), arguing that a social planner behind a veil of ignorance, unsure of the generation she was to be born into, would choose to maximise steady-state welfare.

There is a long tradition in economic policy design that recommends a higher societal discount factor relative to private-sector preferences. Ramsey (1928) provided a famous early articulation of this, deliberately eschewing discounting in his presentation of the consumption-savings problem. 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 far more substantial change to the principles of policy design than our paper. The method that we outline above is deliberately designed to preserve standard solution techniques in time-consistent environments. A change to the societal discount factor will not do this. It will imply intervention to boost the savings rate in efficient economies, simply because the social preference puts greater weight on future generations than the private sector. As we show in the examples, immiseration is not a property of time-consistently undominated policy. Thus if the motivation of Phelan (2006) and Farhi and Werning (2007) is simply to address this particular outcome of Ramsey choice, time-consistently undominated policy may offer a more satisfactory approach.

[60] This was shown independently by Green (1987) and Thomas and Worrall (1990).

[61] Kocherlakota (2010), Chapter 3, argues that the generality of the result should not be exaggerated, providing examples of preferences in a dynamic Mirrlees problem where it does not apply.

[62] For instance, an Atkeson-Lucas environment with log consumption utility and $\beta^{-1} = R$ would see expected consumption remain constant through time.

[63] Farhi and Werning (2010) apply their 2007 methodology to the design of estate taxation.

11.3.2 Positive literature

Closely related to these papers is the small positive literature on 'credible' social insurance policies in environments without commitment. Sleet and Yeltekin (2006) study a dynamic hidden information model in the spirit of Atkeson and Lucas (1992). They investigate optimal policies for the initial period's policymaker, under the assumption that each period's policymaker must be made at least as well off as under the worst available subgame-perfect continuation equilibrium. Their main result is that this is observationally equivalent to a setting with full commitment and a social planner with a higher discount factor than the private sector. Intuitively, the threat of a subsequent deviation provides the initial planner with an incentive to 'backload' welfare.[64]

Clearly this work differs significantly from ours by not assuming a commitment device. More subtly, it also retains the basic 'Ramsey' perspective that, subject to sustainability constraints, a chosen policy should be one that is best under the first period's welfare calculation. By contrast, our paper starts from the perspective that any choice criterion must be symmetrically applicable through time. Choice is then endowed with desirability properties that are only as strong as time-consistent applicability will allow. Thus Sleet and Yeltekin (2006) still recommend a policy that is different in character at the start of time by comparison with subsequent periods, whereas we do not.

A closer positive link is to the work of Kocherlakota (1996b), who introduces the refinement of *reconsideration-proofness* to a game-theoretic setting. This recommends finding the best sustainable equilibrium in a stationary environment, subject to the assumption that future policymakers will be allowed to select across equilibria the same manner. This naturally leads to the best *constant* choice over time, given the assumed absence of state variables. Applying Kocherlakota's refinement to the Atkeson-Lucas problem, Sleet and Yeltekin (2006) indeed show that it is equivalent to maximising steady-state welfare – i.e., letting the discount factor approach unity. This exactly coincides with our time-consistently undominated policy in this example.

With state variables, however, the picture is more complex.[65] Time-consistently undominated policy cannot be recovered by a simple change to the discount factor, and nor can credible social insurance policy – with or without the reconsideration-proofness refinement. For instance, Farhi, Sleet, Werning and Yeltekin (2012) study credible policies in a dynamic Mirrlees economy with capital, and show that although capital accumulation will be affected by a strategic incentive to influence future sustainability constraints, this will not generally correspond to a simple increase in β.[66] Thus dynamic credibility arguments

[64] Related results can be found in Sleet and Yeltekin (2008), Acemoglu, Golosov and Tsyvinski (2010), Farhi, Sleet, Werning and Yeltekin (2012) and Golosov and Iovino (2014).

[65] Both Kocherlakota (1996b) and Sleet and Yeltekin (2006) assumed a stationary environment – i.e., no endogenous states.

[66] C.f. their Proposition 6 in particular, and the arguments preceeding it. Aggregate accumulation depends on the derivative of autarky value with respect to the capital stock – an equilibrium object that is likely increasing in capital. This provides an incentive to restrain accumulation – the opposite of what a higher β would imply.

certainly cannot endorse the *universally* higher discount factor of Phelan (2006) and Farhi and Werning (2007).

11.4 Ramsey taxation

The problem of initial-period policy differing substantially from long-run policy has also been encountered in the Ramsey optimal taxation literature, particularly the classic literature on capital taxation, though here there are some important distinctions to draw. In representative-agent capital tax models that allow the government to accumulate assets, it is well known that a first-best allocation can be achieved when initial tax rates are unconstrained.[67] Since the initial capital levy is equivalent to a lump-sum tax, a high enough initial tax rate can provide the government with enough revenue to fund all of its expenditures in perpetuity, obviating the need for any future taxes. This is widely viewed as trivialising the problem, and so it is conventional to impose some restriction to rule it out. The main options are either (a) to place an upper bound on the initial tax rate, (b) to assume the initial tax rate is predetermined, or (c) to assume that the government must run a balanced budget.[68]

Note that this is a different 'initial period' problem from the pure time variation in choice that arises under Ramsey policy. Without some constraints on initial instrument choice, the Ramsey problem in Chamley-Judd settings can become trivial – and fully time-consistent.[69] But even with these restrictions, the Ramsey-optimal capital tax rate may be systematically higher – or lower – in early periods relative to later ones. Our paper is more concerned with avoiding this latter issue – taking as given that governments themselves may wish to avoid it. A lively current literature has highlighted that the way in which the 'trivialisation' issue is overcome has important implications for long-run Ramsey outcomes.[70] However it has been overcome, our question is how a time-consistent normative choice strategy can then be formulated.

This said, the Straub and Werning (2014) results for Ramsey policy can be viewed as strengthening our motivation. In a version of the capital tax problem due to Judd (1985), they show it can be Ramsey-optimal to drive the long-run value of the capital stock to an extremely low 'subsistence' level, with workers' consumption converging to zero. This occurs even when workers' welfare is the sole concern of the policymaker. It happens in specifications where long-run economic stagnation is an effective way to incentivise early savings. Echoing the immiseration results, this highlights that the long-run outcome under Ramsey policy need not exhibit any clear desirability properties per se. A time-consistently undominated policy overcomes this result, since prior incentives are

[67] See Atkeson, Chari and Kehoe (1999) for a useful discussion.

[68] The first approach is taken by Chamley (1986), the second by Atkeson, Chari and Kehoe (1999) and the third by Judd (1985). Straub and Werning (2014) discuss the convergence implications of approaches (a) and (c), showing that they do not imply zero long-run capital taxes as commonly as first thought.

[69] When a first-best is available, the government has no incentive to deviate from it at any point in time.

[70] See, in particular, the paper by Straub and Werning (2014) .

effectively given less weight in long-run policy.

12 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 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 normatively appealing choices that differ from Ramsey policy both in the short run and the long run. We have shown this both in a general setting, and in a number of canonical 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 an uncooperative equilibrium.[71] But the commitment assumption raises a slightly uncomfortable question. If it were indeed possible for the policymaker in period 0 to commit, why would they ever choose a policy that was not optimal for that period?

In our view this is too direct a reading of the problem. In practice no society has access to a perfect commitment device *ex-ante*, resistant to all conceivable insurgencies and revolutions. Laws can always be repealed, and constitutions can always be torn apart. But the likelihood of this is surely not exogenous to the choice procedure that is applied. A commitment that can survive reassessment according to the same principles that generated it in the first place has an important robustness quality. Popular acceptance of these principles *is* a commitment device.

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.

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

[71] There is a clear analogy here with the distinction between cooperative approaches to game theory, in the spirit of Nash (1950) and Shapley (1953), and more 'standard' non-cooperative approaches, following Nash (1951).

[4] Andrés, J., O. Arce and C. Thomas (2013), 'Banking Competition, Collateral Constraints, and Optimal Monetary Policy', *Journal of Money, Credit and Banking*, 45(s2), 87–125.

[5] Armenter, R. (2017), 'The Perils of Nominal Targets', *The Review of Economic Studies*, forthcoming.

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

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

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

[9] Benigno, P., G. Eggertsson and F. Romei (2016), 'Dynamic Debt Deleveraging and Optimal Monetary Policy', CEPR Discussion Paper DP11180

[10] Benigno, P., and L. Paciello (2014), 'Monetary policy, Doubts and Asset Prices', *Journal of Monetary Economics*, 64, 85–98

[11] Benigno, P., and M. Woodford (2005), 'Inflation Stabilization and Welfare: The Case of a Distorted Steady State', *Journal of the European Economic Association*, 3(6), 1185–1236.

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

[13] Bianchi, J., and E.G. Mendoza (2013), 'Optimal Time-Consistent Macroprudential Policy', NBER WP w19704.

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

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

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

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

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

[19] 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.

[20] 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.

[21] Cúrdia, V., and M. Woodford (2016), 'Credit Frictions and Optimal Monetary Policy', *Journal of Monetary Economics*, 84, 30–65.

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

[23] Diamond, P.A., and J.A. Mirrlees (1978), 'A Model of Social Insurance with Variable Retirement', *Journal of Public Economics*, 10(3), 295–336.

[24] 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.

[25] 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.

[26] Engel, C. (2014), 'Exchange Rate Stabilization and Welfare', *Annu. Rev. Econ.*, 6(1), 155–177.

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

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

[29] Farhi, E., C. Sleet, I. Werning and S. Yeltekin (2012), 'Non-linear Capital Taxation Without Commitment', *Review of Economic Studies*, 79(4), 1469–1493.

[30] Giannoni, M.P., and M. Woodford (2017), 'Optimal Target Criteria for Stabilization Policy', *Journal of Economic Theory*, 168, 55–106.

[31] Golosov, M., and L. Iovino (2014), 'Social Insurance, Information Revelation, And Lack Of Commitment', NBER WP w20633.

[32] Golosov, M., N.R. Kocherlakota and A. Tsyvinski (2003), 'Optimal Indirect and Capital Taxation' *The Review of Economic Studies*, 70(3), 569–587.

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

[34] 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.

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

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

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

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

[39] Kocherlakota, N.R. (1996b), 'Reconsideration-Proofness: A Refinement for Infinite Horizon Time Inconsistency', *Games and Economic Behavior*, 15, 33–54.

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

[41] 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.

[42] 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.

[43] Luenberger, D.G. (1969), *Optimization by Vector Space Methods*, New York: Wiley.

[44] Marcet, A., and R. Marimon (1998), 'Recursive Contracts', European University Institute WP ECO 1998 #37.

[45] Marcet, A., and R. Marimon (2017), 'Recursive Contracts', Manuscript.

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

[47] Nash, J.F. (1950), 'The Bargaining Problem', *Econometrica*, 18(2), 155–162.

[48] Nash, J.F. (1951), 'Non-Cooperative Games', *Annals of Mathematics*, 54(2), 286–295.

[49] 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.

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

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

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

[53] Ramsey, F.P. (1927), 'A Contribution to the Theory of Taxation', *The Economic Journal*, 37(145), 47–61.

[54] Ramsey, F.P. (1928), 'A Mathematical Theory of Saving', *The Economic Journal*, 38(152), 543–559.

[55] Rawls, J. (1971), *A Theory of Justice*, Cambridge, MA: Harvard University Press.

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

[57] Shapley, L.S. (1953), 'A Value for n-person Games', *Contributions to the Theory of Games*, 2(28), 307–317.

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

[59] Sleet, C., and S. Yeltekin (2008), 'Politically Credible Social Insurance', *Journal of Monetary Economics*, 55(1), 129–151.

[60] Straub, L., and I. Werning (2014), 'Positive Long Run Capital Taxation: Chamley-Judd Revisited', NBER WP w20441.

[61] 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.

[62] 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.

[63] Werning, I. (2002), 'Optimal Dynamic Taxation', Ph.D. thesis, University of Chicago.

[64] 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.

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

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:

$$y_t^c = \frac{\chi(1-\beta)^2}{\gamma^2+\chi(1-\beta)^2}y^* \tag{34}$$

$$\pi_t^c = \frac{\chi\gamma(1-\beta)}{\gamma^2+\chi(1-\beta)^2}y^* \tag{35}$$

For this policy to be dominated in some period s, there would have to exist an alternative policy $(\bar{\mathbf{y}}'_{\mathbf{s}}, \bar{\pi}'_{\mathbf{s}})$ 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 $(\bar{\mathbf{y}}^{\mathbf{c}}_{\mathbf{s}}, \bar{\pi}^{\mathbf{c}}_{\mathbf{s}})$ along the vector $[(\bar{\mathbf{y}}^{\mathbf{c}}_{\mathbf{s}}, \bar{\pi}^{\mathbf{c}}_{\mathbf{s}}) - (\bar{\mathbf{y}}'_{\mathbf{s}}, \bar{\pi}'_{\mathbf{s}})]$ 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 $(\bar{\mathbf{y}}'_{\mathbf{s}}, \bar{\pi}'_{\mathbf{s}})$ is a strict improvement on $(\bar{\mathbf{y}}^{\mathbf{c}}_{\mathbf{s}}, \bar{\pi}^{\mathbf{c}}_{\mathbf{s}})$ 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 $r \geq s$:

$$\sum_{t=r}^{\infty}\beta^{t-r}\left[\pi_t^c\frac{d\pi_t}{d\Delta} + \chi\left(y_t^c - y^*\right)\frac{dy_t}{d\Delta}\right] \leq -\delta \tag{36}$$

From the Phillips curve constraint, we know:

$$\frac{dy_t}{d\Delta} = \frac{1}{\gamma}\left[\frac{d\pi_t}{d\Delta} - \beta\frac{d\pi_{t+1}}{d\Delta}\right] \tag{37}$$

Substituting this into inequality (36) gives:

$$\frac{-\beta\chi\gamma}{\gamma^2+\chi(1-\beta)^2}y^*\left\{\sum_{t=r}^{\infty}\beta^{t-r}\frac{d\pi_t}{d\Delta} - \sum_{t=r+1}^{\infty}\beta^{t-r-1}\frac{d\pi_t}{d\Delta}\right\} \leq -\delta \tag{38}$$

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 (38) can be rewritten as:

$$D_{r+1} \leq D_r - \tilde{\delta} \tag{39}$$

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 3

Part 1 of the Proposition is trivial. Suppose otherwise, and take an allocation $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in D(x_{s-1})$ that does not solve Problem 1 for the promise values that it induces. Then there is an allocation $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ in the constraint set for Problem 1 with $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \succ (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$. But belonging to the constraint set for Problem 1 implies that $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ satisfies conditions (11) and (12) for the promise values that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ induces. Thus $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$, and so by Condition 1 $(\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \succ^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$. This contradicts $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}}) \in D(x_{s-1})$.

For part 2 it is sufficient to show that every feasible policy that is time-consistently comparable to $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ must belong to the constraint set for Problem 1, given the promise values that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ induces. This is immediate: conditions (11) and (12), that characterise time-consistent comparability, coincide exactly with (13) and (14), defining the constraint set for Problem 1.

Proof of Proposition 4

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

Consider the promise sequence that $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ induces, denoted $\bar{\omega}''_{\mathbf{s}}$. If $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \in \Xi^g\left(x'_{s-1}\right)$, then it is immediate that $\bar{\omega}''_{\mathbf{s}}$ time-consistently dominates $\bar{\omega}'_{\mathbf{s}}$, since a switch to $\bar{\omega}''_{\mathbf{t}}$ can guarantee at least as desirable an outcome as $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ for all $t \geq s$. Thus $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}}) \notin \Xi^g\left(x'_{s-1}\right)$. But then it follows from the definition of irrelevant extensions that there is a set of alternative allocations $(\bar{\mathbf{x}}'''_{\mathbf{t}}, \bar{\mathbf{a}}'''_{\mathbf{t}}) \in \Xi^g\left(x'_{t-1}\right) \cap \Xi^h$ for all $t \geq s$ (with the chosen $(\bar{\mathbf{x}}'''_{\mathbf{t}}, \bar{\mathbf{a}}'''_{\mathbf{t}})$ potentially varying in t) such that $(\bar{\mathbf{x}}'''_{\mathbf{t}}, \bar{\mathbf{a}}'''_{\mathbf{t}}) \succ^{TC} (\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$. This ordering either applies through constraint dominance or preference dominance. For constraint dominance to apply, by Proposition 3 each $(\bar{\mathbf{x}}'''_{\mathbf{t}}, \bar{\mathbf{a}}'''_{\mathbf{t}})$ must satisfy the constraint set for Problem 1 generated by the promise sequence $\bar{\omega}''_{\mathbf{t}}$, and deliver higher welfare than $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ for all $t \geq s$. Since $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ in turn delivers higher welfare than $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ for all $t \geq s$, including at the limit, it follows that $\bar{\omega}''_{\mathbf{t}} \succ^{\omega}_{x'_{t-1}} \bar{\omega}'_{\mathbf{t}}$ for all $t \geq s$, including at the limit – a contradiction.

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

For the 'only if' part, suppose otherwise. Then there is no alternative allocation $(\bar{\mathbf{x}}^*_{\mathbf{s}}, \bar{\mathbf{a}}^*_{\mathbf{s}}) \in \tilde{\Xi}^g\left(x'_{s-1}\right) \cap \Xi^h$ such that $(\bar{\mathbf{x}}^*_{\mathbf{s}}, \bar{\mathbf{a}}^*_{\mathbf{s}}) \succ^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$, where $\tilde{\Xi}^g\left(x'_{s-1}\right)$ is an irrelevant extension of $\Xi^g\left(x'_{s-1}\right)$, but there is an alternative promise sequence $\bar{\omega}''_{\mathbf{s}}$ that time-consistently dominates $\bar{\omega}'_{\mathbf{s}}$ when the initial state vector is x'_{s-1}. Since $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ solves Problem 1 for the promise sequence $\bar{\omega}'_{\mathbf{s}}$, this means that for all $t \geq s$ (and at any limit as $t \to \infty$) there exists a sequence $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}}) \in \Xi^g\left(x'_{t-1}\right) \cap \Xi^h$ such that $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}}) \succ (\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$, with $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ satisfying the constraints for Problem 1 when the promise sequence is $\bar{\omega}''_{\mathbf{t}}$. Denote by W''_t the value of the social welfare criterion when $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ is implemented, and W'_t when $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'_{\mathbf{t}})$ is implemented. Likewise, r'_t is used as shorthand for $\sum_{\sigma_t \in \Sigma} r\left(a'_t\left(\sigma_t\right), \sigma_t\right) \Pi\left(\sigma_t\right)$, equivalently for r''_t, and so on. $\bar{r} := \sup_{a_t \in A} \sum_{\sigma_t \in \Sigma} r\left(a'_t\left(\sigma_t\right), \sigma_t\right) \Pi\left(\sigma_t\right)$ is an upper bound on r_t, whose existence follows from Assumption 2.

By the definition of time-consistent 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 $t \geq s$:

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

$W''_t - W'_t \geq \varepsilon$ 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 3.6, it is always possible to find a sequence $\bar{\mathbf{a}}'''_{\mathbf{s}} \in \mathcal{A}$ such that $\bar{\mathbf{a}}'''_{\mathbf{s}}$ induces the promise sequence $\bar{\omega}''_{\mathbf{s}}$ and implies a value for the policy criterion of r'''_t for all $t \geq s$. Now suppose the constraint set $\Xi^g\left(x'_{s-1}\right)$ is expanded to $\Xi^g\left(x'_{s-1}\right) \cup (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'''_{\mathbf{s}})$, and let W'''_t be the value of the discounted social welfare criterion in period t when $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'''_{\mathbf{s}})$ is implemented. By construction, $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'''_{\mathbf{t}})$ induces a promise sequence that is also satisfied by $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ for all $t \geq s$, and $W'''_t < W''_t$ for all $t \geq s$, so $(\bar{\mathbf{x}}''_{\mathbf{t}}, \bar{\mathbf{a}}''_{\mathbf{t}})$ constraint-dominates $(\bar{\mathbf{x}}'_{\mathbf{t}}, \bar{\mathbf{a}}'''_{\mathbf{t}})$ for all $t \geq s$. Hence $\Xi^g\left(x'_{s-1}\right) \cup (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'''_{\mathbf{s}})$ is an irrelevant extension of $\Xi^g\left(x'_{s-1}\right)$. But $W'''_t - W'_t \geq \frac{\varepsilon}{2} > 0$ for all $t \geq s$, and $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'''_{\mathbf{s}})$ and $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ imply the same state vector in every period, so $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'''_{\mathbf{s}}) \succ^{TC} (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$. Hence $(\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ cannot belong to $D^*\left(x_{s-1}\right)$.

Proof of Proposition 7

Existence of the saddle point multipliers follows from direct application of Theorem 1, §8.3 in Luenberger (1969), given assumptions 2, 4, 5 and 3. The promise-keeping constraint can be rewritten for all t and σ as:

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

where $\gamma_t\left(\sigma\right) := \omega_t\left(\sigma\right) - \beta \mathbb{E}_t\left[\omega_{t+1}\left(\sigma'\right) \middle| \sigma\right]$, so that the vector movement in promises $\bar{\mathbf{w}}_{\mathbf{s}}$ causes a per-unit change in $\gamma_t\left(\sigma\right)$ of $w_t\left(\sigma\right) - \beta \mathbb{E}_t\left[w_{t+1}\left(\sigma'\right) \middle| \sigma\right]$. Hence, applying Theorem 1, §8.5 of Luenberger (1969), where the derivative exists it is given by:

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

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.)

Proof of Proposition 6

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

$$V\left(\alpha\bar{\omega}'_{\mathbf{s}} + (1-\alpha)\,\bar{\omega}''_{\mathbf{s}}; x_{s-1}\right) \geq \alpha V\left(\bar{\omega}'_{\mathbf{s}}; x_{s-1}\right) + (1-\alpha)\, V\left(\bar{\omega}''_{\mathbf{s}}; x_{s-1}\right) \tag{40}$$

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

Let $\bar{\mathbf{y}}' := (\bar{\mathbf{x}}'_{\mathbf{s}}, \bar{\mathbf{a}}'_{\mathbf{s}})$ and $\bar{\mathbf{y}}'' := (\bar{\mathbf{x}}''_{\mathbf{s}}, \bar{\mathbf{a}}''_{\mathbf{s}})$ solve Problem 1 for $\bar{\omega}'_{\mathbf{s}}$ and $\bar{\omega}''_{\mathbf{s}}$ respectively. It follows from the concavity of r (Assumption 5) that (40) must be satisfied provided the convex combination $\alpha\bar{\mathbf{y}}' + (1-\alpha)\,\bar{\mathbf{y}}''$ is feasible when the promise sequence is $\alpha\bar{\omega}'_{\mathbf{s}} + (1-\alpha)\,\bar{\omega}''_{\mathbf{s}}$. 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 (40), which is then a lower bound on $V\left(\alpha\bar{\omega}'_{\mathbf{s}} + (1-\alpha)\,\bar{\omega}''_{\mathbf{s}}; x_{s-1}\right)$.

The quasiconcavity of p 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 (13) and (14) are also satisfied by the convex combination. Consider (14). For all $t \geq s$,we need:

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

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

$$\alpha h\left(a'_t\right) + (1-\alpha)\, h\left(a''_t\right) + \beta\left[\alpha\omega_{t+1} + (1-\alpha)\,\omega_{t+1}\right] \quad \geq \quad \alpha h^0\left(a'_t\right) + (1-\alpha)\, h^0\left(a'_t\right)$$

But by concavity of h:

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

and by convexity of h^0:

$$h^0\left(\alpha a_t' + (1-\alpha) a_t''\right) \quad \leq \quad \alpha h^0\left(a_t'\right) + (1-\alpha) h^0\left(a_t'\right)$$

Collecting together, this establishes the desired inequality. An identical argument confirms that (13) 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 $\alpha \omega_{\mathbf{s}}' + (1-\alpha) \omega_{\mathbf{s}}''$, completing the first part of the proof. The second part follows by identical logic, noting additionally in this case that strict concavity of r implies $\alpha \bar{\mathbf{y}}' + (1-\alpha) \bar{\mathbf{y}}''$ must deliver *strictly* higher welfare than $\min\left\{V\left(\bar{\omega}_{\mathbf{s}}'; x_{s-1}\right), V\left(\bar{\omega}_{\mathbf{s}}''; x_{s-1}\right)\right\}$ whenever $V\left(\bar{\omega}_{\mathbf{s}}'; x_{s-1}\right) \neq V\left(\bar{\omega}_{\mathbf{s}}''; x_{s-1}\right)$.

Proof of Proposition 8

Let $\bar{\omega}_{\mathbf{t}}'$ be the promise sequence induced by the Ramsey allocation, and consider the directional derivative $\delta_V\left(\bar{\omega}_{\mathbf{t}}', x_{t-1}'; \bar{\mathbf{w}}_{\mathbf{t}}\right)$ for some $t > s$. Rearranging the result in Proposition 7, this derivative will be given by:

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

where $w_t(\sigma) \in \mathbb{R}^j$ denotes the component of $\bar{\mathbf{w}}_{\mathbf{s}}$ 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 (16) to simplify. The result follows by noting that any vector of derivatives $\bar{\mathbf{w}}_{\mathbf{s}}$ 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 $V\left(\bar{\omega}_{\mathbf{t}}', x_{t-1}'\right)$ for all $t > s$, bounded above zero. Thus by Proposition 4, $\left(\bar{\mathbf{x}}_{\mathbf{t}}', \bar{\mathbf{a}}_{\mathbf{t}}'\right)$ cannot belong to $D\left(x_{t-1}'\right)$ for any $t \geq s$.

Proof of Proposition 9

By Proposition 3, if $\left(\bar{\mathbf{x}}_{\mathbf{t}}', \bar{\mathbf{a}}_{\mathbf{t}}'\right)$ belongs to $D\left(x_{t-1}'\right)$ for all $t \geq s$, $\left(\bar{\mathbf{x}}_{\mathbf{t}}', \bar{\mathbf{a}}_{\mathbf{t}}'\right)$ must solve Problem 1 for the promises that this allocation induces, denoted $\omega_{\mathbf{t}}'$. Thus by Proposition 4 it must be the case that $\omega_{\mathbf{t}}'$ 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 $\omega_{\mathbf{t}}'$ must be strictly interior to

$\Omega(x_{s-1})$. As above, let $\delta_V\left(\bar{\omega}'_{\mathbf{s}}, x'_{s-1}; \bar{\mathbf{w}}_{\mathbf{s}}\right)$ be the directional (Gateaux) derivative of $V\left(\bar{\omega}'_{\mathbf{s}}, x'_{s-1}\right)$ as $\bar{\omega}_{\mathbf{s}}$ is varied along dimension $\bar{\mathbf{w}}_{\mathbf{s}}$, and note that $\bar{\mathbf{w}}_{\mathbf{s}}$ is required to be an element of the same vector space as $\bar{\omega}'_{\mathbf{s}}$ (with $w_t(\sigma) \in \mathbb{R}^j$ denoting the component of $\bar{\mathbf{w}}_{\mathbf{s}}$ 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 7 shows, the derivative at differentiable points can be written as:

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

Now fix some $\sigma \in \Sigma$, and suppose that there is a period τ such that the terms $\left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right]$ 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:

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

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 arbirtary constant $\varepsilon_t > 0$ by choosing $w_{t+1}(\sigma)$ to satisfy:

$$\beta\left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right] w_{t+1}(\sigma) - \lambda_t^k(\sigma) w_t(\sigma) \geq \varepsilon_t \tag{41}$$

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 (41) 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 $\omega'_{\mathbf{t}}$ is undominated according to the ordering $\succ^{\omega}_{x'_{t-1}}$ for all $t \geq \tau$.

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

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

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

$$\frac{w_{t+1}(\sigma)}{w_t(\sigma)} \geq (1+\gamma) \frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right]} \tag{42}$$

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:

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

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:

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

or:

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

Using this in (42) confirms that (41) 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 $K\left(\frac{1}{\rho}\right)^{T-\tau} < \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}$ 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:

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

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:

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

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

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

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 $\omega'_{\mathbf{t}}$ is undominated according to the ordering $\succ^{\omega}_{x'_{t-1}}$ 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 $\bar{\omega}'_{\tau}$.[72] Again, if this sequence of differential changes is such that inequality (41) 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 $\omega'_{\mathbf{t}}$ 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 $\prod_{t=\tau}^{T-1}\left|\frac{\omega_t(\sigma)}{\omega_{t+1}(\sigma)}\right|\frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}$ to zero, i.e. for any $\tau \geq s$ there exists a $\rho \in (0,1)$ and $K > 0$ such that for all $T > \tau$:

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

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

$$\left|\frac{\tilde{w}_{t+1}(\sigma)}{\tilde{w}_t(\sigma)}\right| \geq (1+\gamma)\frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}\left|\frac{\omega_t(\sigma)}{\omega_{t+1}(\sigma)}\right| \tag{43}$$

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

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

and the object on the left-hand side converges to zero, whilst the possibility of the lower bound is trivial. Note that if $sign(\omega_t(\sigma)) = sign(\omega_{t+1}(\sigma))$, condition (43) simply states:

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

[72] Part 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.

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

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

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

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

or:

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

Using this in (43), and multiplying through by $\beta\left[\lambda_t^m\left(\sigma\right)+\lambda_t^k\left(\sigma_-\right)\right]\omega_{t+1}\left(\sigma\right)\tilde{w}_t\left(\sigma\right)$, confirms that (41) is satisfied,[73] and the bounds on $\lambda_t^k\left(\sigma\right)\omega_t\left(\sigma\right)$ and $\tilde{w}_t\left(\sigma\right)$ imply ε_t is bounded above zero as required.

The case where $K\left(\frac{1}{\rho}\right)^{T-\tau} < \left|\frac{\omega_\tau(\sigma)}{\omega_T(\sigma)}\right| \prod_{t=\tau}^{T-1} \frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]}$ 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 10

The proof adopts the same approach as for Proposition 9, showing that a differential change to promises can generate an improvement for all policymakers when the stated conditions are not met. Suppose indeed that they are not. Then under **difference comparability** there exists 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:

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

for some $\varepsilon > 0$. Now consider the differential change to $\bar{\omega}_{\mathbf{s}}$ given by $\bar{\mathbf{w}}_{\mathbf{s}}$ such that $w_{t+\tau}\left(\sigma\right) = -1$ 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:

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

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

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

[73] Note that the sign of this expression will be negative if and only if $sign\left(\omega_t\left(\sigma\right)\right) \neq sign\left(\omega_{t+1}\left(\sigma\right)\right)$.

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 $\beta^{T-1}\varepsilon$. 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 $\omega_{\mathbf{s}}$ is dominated. A near-identical argument applies under **ratio comparability**, allowing for the fact that Fréchet derivatives can only be taken as the limit of proportional changes to the promise values in this case.

Proof of Proposition 11

The fact that this cannot be true for $\alpha_t > 1$ follows directly from Proposition 10, and $\alpha_t \geq 0$ follows from the fact that the multipliers are weakly positive. Now 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 implies that there must be at least one degree of linear independence across (positive-measure values of) σ between the values of $\left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right]$ and of $\lambda_{t+1}^k(\sigma)$. Hence, under **difference comparability**, it is possible to find bounded differential changes $\{w_{t+1}(\sigma)\}_{\sigma\in\Sigma}$ such that the following two restrictions are met:

$$\int_{\sigma\in\Sigma} \lambda_{t+1}^k(\sigma)\, w_{t+1}(\sigma)\, d\Pi(\sigma) = 0 \tag{44}$$

$$\int_{\sigma\in\Sigma} \left[\lambda_t^m(\sigma) + \lambda_t^k(\sigma_-)\right] w_{t+1}(\sigma)\, d\Pi(\sigma) = 1 \tag{45}$$

Now, if the condition in the proof is not satisfied, then there is a positive-measure subset of Σ that violates condition (22) 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 these preceeding conditions (44) and (45) with values for $w_t(\sigma)$ that are bounded below some $\bar{w}$, uniform in t. 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 two propositions. This contradicts that the original policy was time-consistently undominated, given Proposition 4. The case of **ratio comparability** again proceeds on the same lines, taking the Gateaux derivative as the limit of a proportional change to each promise value.

Proof of Proposition 12

Suppose the claim is not true. Then under **difference comparability** there exist time periods t such that for a positive-measure subset of σ, the following are true:

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

Following the logic of the previous three Propositions, consider a differential change $w_{t+1}(\sigma) = 1$, applied in all such date-states. The marginal value of this is at least $\beta\varepsilon$ 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 there is a strict 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.

Proof of Proposition 13

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 4, it is sufficient to show that when the three specified conditions are satisfied, there is no marginal change to the promises $\bar{\mathbf{w}}_{\mathbf{s}}$ such that $\delta\left(V\left(\bar{\omega}'_{\mathbf{t}}, x'_{t-1}\right), \bar{\mathbf{w}}_{\mathbf{t}}\right)$ will be bounded above zero for all t sufficiently large, including at the limit as $t \to \infty$.

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

$$\eta_s : \quad = \quad 1$$

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(a):

$$\frac{\lambda^k_{t+1}(\sigma)\left(1 + \Delta_t(\sigma)\right)}{\alpha_t\left[\lambda^m_t(\sigma) + \lambda^k_t(\sigma_-)\right]} \equiv 1 \tag{46}$$

Note that linear convergence implies that the product:

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

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

Condition 2(b) in the Proposition states that the following object is bounded in r for all $\sigma \in \Sigma$:

$$\prod_{t=\tau}^{r-1}\frac{\lambda^k_t(\sigma)}{\beta\left[\lambda^m_t(\sigma) + \lambda^k_t(\sigma_-)\right]}$$

Applying the identity (46), this can be rewritten as follows:

$$\begin{aligned}
\prod_{t=\tau}^{r-1} \frac{\lambda_t^k(\sigma)}{\beta\left[\lambda_t^m(\sigma)+\lambda_t^k(\sigma_-)\right]} &= \prod_{t=\tau}^{r-1} \frac{\alpha_t}{\beta} \frac{\lambda_t^k(\sigma)}{\lambda_{t+1}^k(\sigma)(1+\Delta_t(\sigma))} \\
&= \frac{\lambda_\tau^k(\sigma)}{\lambda_r^k(\sigma)} \prod_{t=\tau}^{r-1} \frac{\eta_{t+1}}{\eta_t(1+\Delta_t(\sigma))} \\
&= \frac{\lambda_\tau^k(\sigma)}{\lambda_r^k(\sigma)} \frac{\eta_r}{\eta_\tau} \prod_{t=\tau}^{r-1} \frac{1}{(1+\Delta_t(\sigma))} \qquad (47)
\end{aligned}$$

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 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 in t for all σ, even if $\sup_t \left\{\lambda_t^k(\sigma)\right\} = \sup_t \{\eta_t\} = \infty$.

Now suppose, contrary to the claim in the Proposition, that there exists a derivative vector $\bar{\mathbf{w}}_{\mathbf{s}}$ that delivers a strict improvement for the policymaker for all $t \geq \tau$, for some $\tau \geq s$. This implies that there must be a scalar $\varepsilon > 0$ such that $\delta_V\left(\bar{\omega}_{\mathbf{t}}', x_{t-1}'; \bar{\mathbf{w}}_{\mathbf{t}}\right) \geq \varepsilon \eta_t$ for all $t \geq \tau$.[74] We will establish a contradiction. As shown in Proposition 7, the Gateaux derivative in all periods $t \geq \tau$ then satisfies:

$$\begin{aligned}
\delta_V\left(\bar{\omega}_{\mathbf{t}}', x_{t-1}'; \bar{\mathbf{w}}_{\mathbf{t}}\right) &= \sum_{r=t}^{\infty} \beta^{r-t} \int_{\sigma\in\Sigma} \left\{\beta\left[\lambda_r^m(\sigma)+\lambda_r^k(\sigma_-)\right] w_{r+1}(\sigma) - \lambda_r^k(\sigma) w_r(\sigma)\right\} d\Pi(\sigma) \\
&\geq \varepsilon \eta_t
\end{aligned}$$

where $w_r(\sigma)$ is the marginal increase in the date-state-specific promise $\omega_r(\sigma)$. Condition 2(a) in the Proposition implies that this inequality can be rewritten as:

$$\sum_{r=t}^{\infty} \beta^{r-t} \int_{\sigma\in\Sigma} \left\{\beta \lambda_{r+1}^k(\sigma) \frac{1}{\alpha_r}\left[1+\Delta_r(\sigma)\right] w_{r+1}(\sigma) - \lambda_r^k(\sigma) w_r(\sigma)\right\} d\Pi(\sigma) \geq \varepsilon \eta_t$$

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

[74] The normalisation by η_t here corrects for the possibility that marginal benefits from increasing promises could be growing without bound in t, in which case vanishingly small changes to ω_t could deliver strictly positive improvements. This would not directly imply the existence of a change to $\bar{\omega}_{\mathbf{t}}$ that is strictly preferred at the limit as $t \to \infty$, since the derivative change to promises at the limit would then be zero. Intuitively, we are requiring that a strict improvement must still exist at the limit when the units in which W_t is expressed are normalised relative to η_t.

Rewriting the previous condition gives:

$$
\begin{aligned}
& \sum_{r=t+1}^{\infty} \beta^{r-(t+1)} \left(\frac{\beta}{\alpha_{r-1}} - \beta \right) \int_{\sigma \in \Sigma} \lambda_r^k (\sigma) \, w_r (\sigma) \, d\Pi (\sigma) \\
\geq \; & \int_{\sigma \in \Sigma} \lambda_t^k (\sigma) \, w_t (\sigma) \, d\Pi (\sigma) + \varepsilon \eta_t \\
& - \sum_{r=t}^{\infty} \beta^{r-t} \int_{\sigma \in \Sigma} \frac{\beta}{\alpha_r} \Delta_r (\sigma) \, \lambda_{r+1}^k (\sigma) \, w_{r+1} (\sigma) \, d\Pi (\sigma)
\end{aligned} \tag{48}
$$

Dividing through by η_t and applying the definition of $\tilde{\lambda}_t^k (\sigma)$ to this gives:

$$
\begin{aligned}
& \sum_{r=t+1}^{\infty} \beta^{r-(t+1)} \left(\frac{\beta}{\alpha_{r-1}} - \beta \right) \frac{\eta_r}{\eta_t} \int_{\sigma \in \Sigma} \tilde{\lambda}_r^k (\sigma) \, w_r (\sigma) \, d\Pi (\sigma) \\
\geq \; & \int_{\sigma \in \Sigma} \tilde{\lambda}_t^k (\sigma) \, w_t (\sigma) \, d\Pi (\sigma) + \varepsilon \\
& - \sum_{r=t}^{\infty} \beta^{r-t} \int_{\sigma \in \Sigma} \frac{\eta_r}{\eta_t} \Delta_r (\sigma) \, \tilde{\lambda}_{r+1}^k (\sigma) \, w_{r+1} (\sigma) \, d\Pi (\sigma)
\end{aligned} \tag{49}
$$

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 (49) 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 (49) 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 a Fréchet derivative to be well defined in the chosen vector space, and the latter was established above under the conditions assumed. Hence we have a contradiction.

The proof under **ratio comparability** proceeds near-identically, allowing for the fact that Fréchet 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.