2-2012

Revision Proofness

Laurence Ales
Carnegie Mellon University, ales@andrew.cmu.edu

Christopher Sleet
Carnegie Mellon University, csleet@andrew.cmu.edu

Revision Proofness

Laurence Ales and Christopher Sleet*

This version: February 2012.

Abstract

We analyze an equilibrium concept called *revision-proofness* for infinite-horizon games played by a dynasty of players. Revision-proofness requires strategies to be robust to joint deviations by multiple players and is a refinement of sub-game perfection. Sub-game perfect paths that can only be sustained by reversion to paths with payoffs below those of an alternative path are not revision-proof. However, for the important class of quasi-recursive games careful construction of off-equilibrium play can render many, and in some cases all, sub-game perfect paths revision-proof.

1 Introduction

We consider infinite horizon repeated and dynamic games played by a sequence or dynasty of strategic players. These players may be interpreted as distinct individuals or as the selves of a single individual with time inconsistent preferences. Following Kocherlakota (1996), they may be interpreted as policymakers in a reduced form representation of a macroeconomic policy game. Infinite horizon dynastic games typically admit large sets of sub-game perfect equilibria. This lack of determinacy coupled with the unappealingly severe nature of some equilibria has motivated a search for refinements. In this paper we explore one such refinement which we label *revision-proofness*. A strategy is revision-proof if there is no alternative strategy that weakly raises the payoffs of all players in a sub-game and strictly raises the payoffs of some. Variations on this definition have been provided by Hammond (1975), Asheim (1997) (from whom we borrow the name revision-proofness) and Caplin and Leahy (2006).[1] However, as yet limited characterization of the concept has been given. Our goal is to fill this gap.

*Tepper School of Business, Carnegie Mellon University, Pittsburgh PA 15217

[1] Alternative renegotiation-proof concepts for repeated games have been provided by, inter alia, Kocherlakota (1996), Pearce (1987) and Farrell and Maskin (1989). We discuss these briefly below.

Sub-game perfection in dynastic games requires that a strategy be robust to a rather limited set of "revisions", namely those that involve a single player altering her action and taking the response of her successors to this alteration as given. Since players do not fully internalize the effect of their actions on other players' payoffs, sub-game perfection often permits strategies that all players agree are weakly inferior and some think are strictly inferior to alternatives. The latter, however, can only be reached via coordinated reforms and these are not possible.

Revision-proofness requires that a strategy be robust to a much larger set of alternatives. Now, all feasible paths of actions are candidate revisions. However, revision-proofness permits a path to disrupt a strategy only if *every* player along the path weakly prefers continuing with it to returning to the strategy and at least one player strictly prefers this. A difficulty with the concept is that play prescribed arbitrarily far into the future is relevant for determining whether or not a strategy is revision-proof. On the one hand, revision-proofness requires that a strategy be robust to coordinated and, given histories, mutually beneficial deviations of arbitrary length; on the other, it permits players arbitrarily far in the future to block a revision that weakly benefits all of their predecessors. As a result, recursive methods are of limited use in characterizing revision-proofness, significantly complicating the analysis.

The extent to which revision-proofness refines sub-game perfection, depends upon the structure of player preferences. We show that if (i) a player strictly prefers a revision path A to a prescribed path B and (ii) deterring each player from being the last to join A "traps" her successor's payoff below that obtained from A, then B cannot be implemented by any revision-proof strategy. We say that the path B is not revision-proof. We give two implications of this result. First, if in every sub-game there is a path whose continuation is optimal for the current and later players, then a strategy is revision-proof if (and only if) it attains such an optimum in every sub-game. Sub-game perfection does not ensure this.[2] Second, if a sub-game perfect strategy uses an exploding sequence of punishments to deter players from unilaterally joining a revision, then it is not revision-proof. If sub-game perfect implementation of a path requires such a sequence, then the path is not revision-proof.

By definition, if a strategy builds up player payoffs along any revision-path so that eventually a player prefers to revert to the strategy, then it is revision-proof. We pursue this logic in the context of quasi-recursive games, an important class of dynastic game. In these games a player and her successor may have differing preferences over

[2] It does so under the stronger condition that player preferences are time consistent and satisfy a "continuity-at-infinity" condition.

the successor's choices, but conditional on the successor's choice, identical preferences over future play. We give examples including a pension game of Hammond (1975), a macro-policy game of Kocherlakota (1996) and a savings game with quasi-hyperbolic discounting. We provide conditions that ensure many, and in some cases all, action paths in quasi-recursive games are revision-proof. To the extent that the conditions we give are satisfied, revision-proofness is a weak refinement with respect to paths and outcomes. On the other hand, it remains restrictive with respect to off-equilibrium play. Our procedure for showing that a given path is revision-proof (i.e. is implemented by at least one revision-proof strategy) is constructive. We build a strategy that implements the path and confronts all sufficiently long revisions with a "blocking" player. For this player there is no alternative continuation path that raises her predecessors' payoffs without reducing her own or her successors' payoffs. Consequently, the blocking player (or her successors) will reject any attempt to revise the strategy towards a continuation path that her predecessors prefer; any successful revision must leave the path following the blocking player intact. The strategy is then constructed to ensure that any revision which does this makes some prior player worse off and, consequently, is revision proof.

The paper proceeds as follows. After a brief review of the literature, a general dynastic game is outlined in Section 2. The formulation allows players to care about past histories of actions in arbitrary ways and so accommodates both repeated games (in which players do not care about the past) and dynamic games (in which they care about a state variable inherited from the past). Section 3 defines sub-game perfection and revision-proofness respectively and introduces some preliminary results. Section 4 provides conditions for a path of actions not to be revision-proof and considers the implications of these conditions for games featuring weak agreement over optima and (potentially) exploding punishment paths. Sections 5 and 6 give conditions for revision-proofness in quasi-recursive games first without and then with state variables. Section 7 relates revision-proofness to Asheim's revision-proofness concept and Kocherlakota's reconsideration-proof concept. Section 8 concludes. Appendices give proofs and supplementary results.

Literature In Hammond (1975) a strategy is defined to be a *dynamic equilibrium* if there exists no alternative strategy that strictly raises the payoffs of players whenever it prescribes different actions for those players. Thus, Hammond's refinement is weaker than that proposed here since we require the absence of an alternative that strictly raises the payoffs of some players in a sub-game, while leaving the payoffs, but not necessarily the actions of the others unchanged. Asheim (1997) proposes a refinement which he labels revision-proofness. We show that although formulated differently the concepts are identical. Asheim (1997) relates revision-proofness to the coalition-proof Nash concept of Bernheim, Peleg, and Whinston (1987) and to the consistent planning concept of Strotz (1956). It is well known that consistent plans need not exist. By extension, revision-proof strategies need not exist. Caplin and Leahy (2006) address the existence problem in finite horizon settings.

Kocherlakota (1996) gives a refinement for dynastic games which he calls *reconsideration-proofness*. A strategy is reconsideration-proof if it is best amongst payoff stationary sub-game perfect strategies. Revision-proofness allows a player to compare two strategies across histories and coordinate her successors onto the dominating one. In contrast, reconsideration-proofness supposes that a player compares the continuations of a given strategy and selects the dominant one subject to the constraint that her successors will do likewise. In the absence of state variables or history dependence, only payoff stationary equilibria survive this internal reconsideration process. Kocherlakota (1996) further assumes that players coordinate onto a best payoff stationary equilibrium. Reconsideration-proofness can be interpreted as a specialization of the renegotiation concept of Farrell and Maskin (1989) to dynastic games.[3] Kocherlakota (1996) develops reconsideration-proofness only for repeated games, not dynamic ones which is an important limitation. To date no extension of the concept to dynamic games has been provided.

Pearce (1987) supplies an alternative notion of renegotiation-proofness for repeated games played by a finite number of infinitely-lived players. He assumes that a strategy will be revised if there is a history and an alternative strategy that gives higher payoffs at all subsequent histories relative to those obtained under the original strategy and at the original history. Again a difficulty with this renegotiation-proof concept is that its extension to dynamic games is unclear.

2 The Environment

We consider games played by a sequence or dynasty of one period-lived players. Players may be interpreted as different individuals or different selves of the same individual. Let $(\mathscr{P}, \rho)$ denote a metric space of actions and $\mathscr{N}$ the natural numbers. Successive players,

[3]Despite the differences, our results showing the limits of revision-proofness as a refinement of sub-game perfection in quasi-recursive games, is in the spirit of van Damme (1989) who shows that Farrell and Maskin (1989)'s concept has no bite in the repeated prisoner dilemma game when discount factors are high enough.

$t \in \mathscr{N}$, choose actions p_t from $\mathscr{P}$. Let $P = \{p_t\}_{t=1}^{\infty}$ be the complete *path* of actions chosen by agents and for $T \in \mathscr{N}$, let $P^T = \{p_t\}_{t=1}^{T}$ be a T-period *history* of actions and $P_{T+1} = \{p_t\}_{t=T+1}^{\infty}$ a period $T+1$ *continuation path*. The sets $\mathscr{P}^{\infty} = \mathscr{P} \times \mathscr{P} \times \ldots$ and $\mathscr{P}^T := \mathscr{P} \times \ldots \times \mathscr{P}$, containing paths and T-period histories respectively, are endowed with the associated product topologies. $\mathscr{P}^0$ is set equal to $\emptyset$ and p_0 and P^0 to null elements. If $T' \geq T$ and the first T elements of $P^{T'}$ equal P^T, then $P^{T'}$ is said to be a successor history to P^T.

The objectives of the players are given by a payoff function $U : \mathscr{N} \times \mathscr{P}^{\infty} \to \overline{\mathbb{R}}$, $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty\}$, with $U(t, \cdot)$ the objective of the t-th player. $U(t, \cdot)$ is defined over the *entire* action path so that the t-th player's payoffs depend upon her ancestors', her own and her descendants' actions. Depending on the setting, this dependence can occur because of preference interactions across individuals (e.g. altruism or envy), or selves (e.g. intra-personal discounting, habit formation). In addition, earlier players may affect the choice sets of later ones by, for example, the accumulation of capital. To economize on notation we encode such effects into player objectives by assigning $-\infty$ payoff values to paths that past play renders infeasible. This becomes a further channel via which players can affect the payoffs of successors.

Definition 1. Given U, $t \in \mathscr{N}$ and $P^{t-1} \in \mathscr{P}^{t-1}$, define the set of *feasible continuation paths* for the t-th player at P^{t-1} by $\Pi_t(P^{t-1}) = \{P \in \mathscr{P}^{\infty} | U(t, P^{t-1}, P) > -\infty\}$. $P = (P^{t-1}, P_t)$ is said to be feasible for the t-th player if $P_t \in \Pi_t(P^{t-1})$. Denote the set of *feasible actions* for the t-th player at P^{t-1} by $\Gamma_t(P^{t-1}) = \{p | \exists P \text{ with } (p, P) \in \Pi_t(P^{t-1})\}$ (i.e. the projection of $\Pi_t(P^{t-1})$ onto its first coordinate).

We impose the following condition.

Assumption 1. *Let U satisfy:*

(i) for all $t \in \mathscr{N}$ and $P^{t-1} \in \mathscr{P}^{t-1}$, $\Pi_t(P^{t-1}) \neq \emptyset$;

(ii) for all $t \in \mathscr{N}$ and $P^{t-1} \in \mathscr{P}^{t-1}$,

$$\Pi_t(P^{t-1}) = \{(p, P) | p \in \Gamma_t(P^{t-1}), P \in \Pi_{t+1}(P^{t-1}, p)\}.$$

Assumption 1 (i) ensures that all players after all histories have a feasible continuation path. Assumption 1 (ii) ensures that if a path is feasible for a player, then it is feasible for all of the player's successors. Conversely, a feasible action choice for the current player combined with any feasible continuation path for the next player forms a feasible continuation path for the current player, (i.e. following a feasible current action choice, a

$-\infty$ payoff cannot be delivered to the current player without the successor incurring such a payoff as well). This formulation is quite general and allows us to derive a number of basic results without imposing much structure on players' objectives. Sections 5 and 6 consider quasi-recursive settings in which player payoffs do not depend on the past or do so through a low dimensional state variable.

Together the set $\mathscr{P}$, sequence of players $\mathscr{N}$ and the payoff function U defines a dynastic game, $\mathscr{G}(U)$. The notation emphasizes the dependence of the game on player objectives. Modulo a relabeling of players each sub-game of $\mathscr{G}(U)$ is itself a dynastic game. The sub-game following P^t will be denoted $\mathscr{G}(U; P^t)$.

A strategy $\sigma = \{\sigma_t\}_{t=1}^{\infty}$, $\sigma_t : \mathscr{P}^{t-1} \to \mathscr{P}$, describes player behavior after every history. Let $(\sigma|P^t) = \{\sigma_{t+r}(P^t, \cdot)\}_{r=1}^{\infty}$ denote the continuation of σ after some history P^t and let $\Phi(\sigma)$ denote the action path induced by strategy σ. We require strategies to be feasible for players in the following sense, for all t and P^{t-1}:[4]

$$U(t, P^{t-1}, \Phi(\sigma|P^{t-1})) > -\infty. \tag{1}$$

Let $\mathscr{S}(U)$ denote the set of feasible strategies for $\mathscr{G}(U)$. Note that Assumption 1 implies that $\mathscr{S}(U) \neq \emptyset$. Of course, if $\sigma \in \mathscr{S}(U)$, then $(\sigma|P^t) \in \mathscr{S}(U; P^t)$. In the remainder of the paper we will use the term strategy to refer to a feasible strategy.

The function U implies indirect payoff functions over players, histories and strategies which we denote, $V_t : \cup_s \mathscr{P}^s \times \mathscr{S} \to \mathbb{R}$ with:

$$\forall P^r \in \cup_s \mathscr{P}^s, \quad V_t(P^r, \sigma) := U(t, P^r, \Phi(\sigma|P^r)).$$

$V_t(P^r, \sigma)$ gives the t-th player's evaluation of the history P^r and the continuation path induced by σ after P^r.

3 Sub-game Perfection and Revision-proofness

We now define sub-game perfection and revision-proofness.

Definition 2. Given a game $\mathscr{G}(U)$, $\sigma \in \mathscr{S}(U)$ is *sub-game perfect* if there is no $t \in \mathscr{N}$,

[4]By Assumption 1, the period t feasibility condition in (1) does not restrict the future actions of players at $t + r$, $r \in \mathscr{N}$, beyond requiring that they too are feasible.

history $P^{t-1} \in \mathscr{P}^{t-1}$ and action $p \in \mathscr{P}$ such that:[5]

$$V_t(P^{t-1}, p, \sigma) > V_t(P^{t-1}, \sigma).$$

$P \in \mathscr{P}^\infty$ is a *sub-game perfect path* if $P = \Phi(\sigma)$ for some sub-game perfect strategy σ.

Definition 3. Given a game $\mathscr{G}(U)$, $\sigma \in \mathscr{S}(U)$ is *revision-proof* if there is no $t \in \mathscr{N}$, history $P^{t-1} \in \mathscr{P}^{t-1}$ and alternative strategy σ' with for all $r \in \mathscr{N}$ and $P^r \in \mathscr{P}^r$:

$$V_{t+r}(P^{t-1}, P^r, \sigma') \geq V_{t+r}(P^{t-1}, P^r, \sigma), \tag{2}$$

and for at least one P^r,

$$V_{t+r}(P^{t-1}, P^r, \sigma') > V_{t+r}(P^{t-1}, P^r, \sigma). \tag{3}$$

A path $P \in \mathscr{P}^\infty$ is *revision-proof* if $P = \Phi(\sigma)$ for some revision-proof strategy σ.

Sub-game perfection rules out strictly profitable unilateral deviations; revision-proofness rules out multilateral deviations that are strictly profitable for some and weakly profitable for all players in a sub-game. Consequently, revision-proofness is a refinement of sub-game perfection, a fact we record below.

Proposition 1. *If σ is revision-proof for $\mathscr{G}(U)$, then it is sub-game perfect for $\mathscr{G}(U)$.*

Proof. The contrapositive is proved. If σ is not sub-game perfect for $\mathscr{G}(U)$, then there is some $P^t = (P^{t-1}, p_t)$ such that: $V_t(P^t, \sigma) > V_t(P^{t-1}, \sigma)$. Let σ' equal σ for all histories except P^{t-1}. Set $\sigma'_t(P^{t-1}) = p_t$. Then $V_t(P^{t-1}, \sigma') > V_t(P^{t-1}, \sigma)$ and for all $r \in \mathscr{N}$ and P^r, $V_{t+r}(P^t, P^r, \sigma') = V_{t+r}(P^t, P^r, \sigma)$. Thus, σ is not revision-proof for $\mathscr{G}(U)$. □

The following result shows that revision-proofness satisfies a weak recursive property.

Proposition 2. *σ is revision-proof for $\mathscr{G}(U)$ if and only if (i) for each $p \in \mathscr{P}$, $(\sigma|p)$ is revision-proof for $\mathscr{G}(U;p)$ and (ii) $\nexists$ a p and σ' for $\mathscr{G}(U;p)$ such that $U(1, p, \Phi(\sigma')) > V_1(\sigma)$ and for all $t = 2, 3, \ldots$ and P^{t-2}, $U(t, p, P^{t-2}, \Phi(\sigma'|P^{t-2})) = V_t(p, P^{t-2}, \sigma)$.*

Proof. Necessity of the conditions in the proposition for revision-proofness is an immediate consequence of the definition. For sufficiency, suppose that σ satisfies the conditions

[5] Note that our restriction to feasible strategies eliminates uninteresting cases in which each player takes an action yielding a $-\infty$ payoff to its predecessor.

in the proposition. Condition (i) ensures that there is no P^{t-1}, $t = 2, 3, \ldots$ and σ' satisfying (2) and (3). Hence, any alternative σ' either lowers the payoff of a player at some $t > 1$ or leaves all such players' payoffs unaltered. It only remains to rule out the possibility that a strategy σ' strictly raises the first player's payoff while leaving all subsequent players' payoffs unchanged. This is done by Condition (ii). □

A strategy σ is sub-game perfect for $\mathscr{G}(U)$ if and only if (i) σ_1 is optimal for the initial player given the continuations $(\sigma|p)$, $p \in \mathscr{P}$, and (ii) each continuation $(\sigma|p)$ is sub-game perfect for $\mathscr{G}(U; p)$. Thus, sub-game perfect strategies have a recursive structure. Sections 5 and 6 impose additional "quasi-recursive" structure on player payoffs.[6] Then, following Abreu, Pearce, and Stacchetti (1990), the recursivity of sub-game perfection may be exploited to show that the set (or correspondence) of sub-game perfect payoffs is a fixed point of a time invariant, Bellman-like operator. In addition, if this set is bounded (or the correspondence has a bounded graph), then this set (or correspondence graph) is the largest bounded fixed point of the Bellman-like operator.

Proposition 2 asserts that σ is revision-proof for $\mathscr{G}(U)$ if and only if (i) it gives an optimal payoff to the initial player subject to all later players receiving payoffs no lower than under σ and (ii) each continuation $(\sigma|p)$ is revision-proof for $\mathscr{G}(U; p)$.[7] Condition (i) implies that the continuations $(\sigma|p)$ of a revision-proof strategy satisfy an infinite horizon refinement of revision-proofness: the initial player can break the indifference of all later players in her favor. In general, the infinite horizon reach of the refinement disrupts full recursivity. Under stronger assumptions on player payoffs, in particular, the quasi-recursiveness of later sections, a strategy σ is revision-proof if (i) is weakened to (i′) σ gives an optimal payoff to the initial player subject to her immediate successor receiving a payoff no lower than under σ and (ii) holds as before. Under quasi-recursiveness, if later players have broken future indifferences in their favor (which (ii) ensures), then the only potential indifferences left for the initial player to break are those of her immediate successor. Ales and Sleet (2011) show that in this case, the set of revision-proof payoffs is a fixed point of a modified version of the operator of Abreu, Pearce, and Stacchetti (1990).[8] Nonetheless, a difficulty with revision-proofness remains even if payoffs are

[6]Payoffs are quasi-recursive if a player and her successor have potentially different preferences over the successor's action given the past, but identical preferences over continuation paths given the past *and* the successor's action.

[7]Since each continuation strategy is revision-proof, opportunities for finding an alternative strategy that weakly raises the payoffs of all and strictly raises the payoffs of some future players are exhausted. Thus, (i) can be stated in terms of all future players receiving payoffs equal to those under σ as is done in Proposition 2.

[8]This operator essentially purges all indifferences that can be broken by prior players from strategies.

quasi-recursive. By definition, if σ is not revision-proof, then there is an alternative σ' that delivers a weakly higher payoff to all players in some sub-game $\mathscr{G}(U;P^t)$ and a strictly higher payoff to at least one. However, play of any such σ' for the first T periods of $\mathscr{G}(U;P^t)$ followed by reversion to σ may make the player T periods into $\mathscr{G}(U;P^t)$ worse off no matter how large T is made. In this case, σ is robust to all finite length revisions, but is still not revision proof. The problem is that revision-proofness permits a player arbitrarily far into the future to block a candidate revision. In quasi-recursive settings, this implies that the largest bounded fixed point of the operator identified in Ales and Sleet (2011) gives the set of payoffs from strategies that are "finite-revision-proof", i.e. robust to all finite length revisions, but it may strictly contain the set of revision-proof payoffs.[9] Thus even if player payoffs are quasi-recursive, recursive methods may fail to provide sufficient conditions for payoffs or paths to be generated by revision-proof strategies. This complicates the analysis of revision-proofness and compels us to pursue non-recursive approaches.

We close this section with a useful lemma which gives an elementary, but convenient path-wise characterization of revision-proofness.

Definition 4. Given $\mathscr{G}(U)$, P' is a successful revision path for σ at P^{t-1} if for all $r \in \mathscr{N}$,

$$U(t-1+r, P^{t-1}, P') \geq V_{t+r}(P^{t-1}, P'^r, \sigma), \tag{4}$$

with strict inequality for at least one r

Lemma 1. *Let Assumption 1 hold. σ is revision-proof for $\mathscr{G}(U)$ if only if there is no history P^{t-1} and path P' such that P' is a successful revision path for σ at P^{t-1}.*

Proof. See Appendix A. □

4 Revision-Proofness in dynastic games

In this section, we explore the extent to which revision-proofness can refine the set of strategies and paths. To fix ideas, we begin with two simple examples.

[9]The situation is quite different from sub-game perfection in games played by a finite number of infinite-lived players. In such games, if player payoffs are "continuous at infinity", then an alternative strategy for a player that raises the player's payoff in a sub-game, must also raise the player's payoff if it is played for a sufficiently large number of periods. Thus, a strategy that is robust to all unilateral finite-length deviations is sub-game perfect.

Example 1. Suppose two actions A and B. Players derive utility from their own and their successor's action choices. Player preferences are given by $U(t,P) = u(p_t) + v(p_{t+1})$, where $u(A) = u(B)$ and $v(A) > v(B)$. Thus, each player is indifferent between taking A or B, but strictly prefers that her successor takes action A. All player's agree that continuation of $A^\infty = (A, A, A, \ldots)$ is optimal. The strategy σ^1 with for all t and P^{t-1}, $\sigma_t^1(P^{t-1}) = A$ is sub-game perfect and delivers the maximal payoff of $u(A) + v(A)$ to all players. Clearly, it is also revision-proof.

Consider another sub-game perfect strategy σ^2, with for all t and P^{t-1}, $\sigma_t^2(P^{t-1}) = B$. In this case, each player is taking a maximal action given continuation play. However, σ^2 is not revision-proof. From any history P^{t-1}, the path A^∞ is a successful revision: it raises the payoff to all players $t-1+r$, $r \in \mathcal{N}$, at histories $A^r = (A, A, \ldots, A)$, by breaking the indifference of players in favor of predecessors.

In fact, strategy σ^1 and the slight variation $\sigma^{1\prime}$, in which the very first action is switched to B, but otherwise A is played, are the only revision-proof strategies in this game. All other strategies deliver weakly lower payoffs at all histories and strictly lower payoffs at some. It follows that all revision-proof strategies attain player optima $\overline{U}(t, P^{t-1})$ in all sub-games $\mathcal{G}(U, P^{t-1})$. $\square$

In Example 1, revision-proofness breaks indifferences of later players in favor of earlier ones and, hence, refines the set of sub-game perfect strategies. Caplin and Leahy (2006) emphasize this aspect of the concept in finite horizon games. In infinite horizon games, revision-proofness may additionally exclude strategies in which some or all players strictly prefer to take a sub-optimal current action given their successors' prescribed responses. In these cases, there may be no indifferences to break (and no successful finite revisions), but an infinite revision may still improve all players' payoffs relative to reversion to the strategy. The next example illustrates.

Example 2. Suppose three actions A, B and C. A player derives utility from her own and her successor's action choices. Her preferences are given by: $U(t,P) = u(p_t) + v(p_{t+1})$, where $u(A) > u(B) \geq u(C)$, $\quad v(A) > v(C) > v(B)$, and $u(C) + v(C) > u(A) + v(B)$. All players agree that A is a weakly optimal action in all periods (and a strictly optimal one during their own lifetime). However, they would rather play C and have their successor do likewise than play A and have their successor play B. It is clear that the strategy σ^1 with for all t and P^{t-1}, $\sigma_t^1(P^{t-1}) = A$ is sub-game perfect and delivers the maximal payoff of $u(A) + v(A)$ to all players. Consequently, it is also revision-proof.

A second sub-game perfect strategy is defined by σ^2, with $\sigma_1^2 = C$ and, for $t > 1$,

$$\sigma_t^2(P^{t-1}) = \begin{cases} C & \text{if } p_{t-1} = \sigma_{t-1}^2(P^{t-2}) \\ B & \text{otherwise.} \end{cases}$$

In this case, the threat of a future play of B deters players from choosing their optimal action A. No player would wish to be the last in a revision consisting of a finite number of plays of A followed by reversion to σ^2. Thus, all revision paths $(A^r, \Phi(\sigma^2|A^r))$ are unsuccessful. But A^∞ *is* a successful revision path: it raises the payoff to players at all sub-histories A^r.

σ^1 is the only revision-proof strategy in this game. Any other strategy, σ^2 for example, delivers a weakly lower payoff at all histories and a strictly lower payoff at some. As in the last example, all revision-proof strategies attain player optima in all sub-games. □

A common element in the previous examples is the existence of a "dominating path" $\widetilde{P} = A^\infty$ and a pair of "dominated continuation path sets", $\mathscr{D}_i$, $i = 0, 1$. In Example 2 these are given by $\mathscr{D}_0 = \mathscr{P}^\infty \backslash \{A^\infty\}$ and $\mathscr{D}_1 = \mathscr{P}^\infty$.[10] The triple $(\widetilde{P}, \mathscr{D}_0, \mathscr{D}_1)$ satisfies two properties. First, for any t and P^{t-1}, the t-th player strictly prefers $\widetilde{P}$ to any path in $\mathscr{D}_0$ and players $t + r$, $r \in \mathscr{N}$, weakly prefer the continuation of $\widetilde{P}$ to any path in $\mathscr{D}_1$. Second, to deter the $t - 1 + r$-th player from defecting away from a path in $\mathscr{D}_0$ (if $r = 1$) or $\mathscr{D}_1$ (if $r > 1$) and playing A, a continuation path in $\mathscr{D}_1$ must be played. This second property is rather trivial in the example since $\mathscr{D}_1 = \mathscr{P}^\infty$. However, in more general settings it is not. These properties ensure that no path in $\mathscr{D}_0$ is revision-proof. Essentially, deterring successive players $t + r$ from being the last to join a finite revision $\widetilde{P}^r$ "traps" continuation play into sets that are payoff dominated by $\widetilde{P}$.

The following proposition extends the intuition of the previous paragraph to settings where the set $\mathscr{D}_1$ might change for every successive player. Starting from t and a prior history P^{t-1}, Proposition 3 supposes a dominating path $\widetilde{P}$ and a sequence of dominated continuation path sets, $\mathscr{D}_r$, $r \in \mathscr{N}$. These satisfy two properties analogous to those above. First, each player $t - 1 + r$, $r \in \mathscr{N}$, weakly prefers (strictly for $r = 1$) the continuation of $\widetilde{P}$ to the paths in $\mathscr{D}_r$-th continuation path set. Second, to deter the $t - 1 + r$-th player from defecting away from a path in $\mathscr{D}_r$ and playing $\widetilde{p}_r$, a continuation path in $\mathscr{D}_{r+1}$ must be played. Proposition 3 shows that this ensures no path in $\mathscr{D}_0$ is revision-proof.

Proposition 3. *Let Assumption 1 hold. Suppose that at P^{t-1}, there is a path $\widetilde{P}$ and a family of path sets $\{\mathscr{D}_r\}_{r=0}^\infty$, $\mathscr{D}_r \subseteq \mathscr{P}^\infty$, such that for each $r \in \{0\} \cup \mathscr{N}$ the following conditions hold:*

[10] In Example 1, $\mathscr{D}_0 = \mathscr{P}^\infty \backslash \{A^\infty, (B, A^\infty)\}$.

(i) if $P \in \mathscr{D}_r$, then $U(t+r, P^{t-1}, \widetilde{P}) \geq U(t+r, P^{t-1}, \widetilde{P}^r, P) > -\infty$ with the first inequality strict if $r = 0$;

(ii) if $P \in \mathscr{D}_r$ and $U(t+r, P^{t-1}, \widetilde{P}^r, P) \geq U(t+r, P^{t-1}, \widetilde{P}^{r+1}, P')$, then $P' \in \mathscr{D}_{r+1}$.

No path in $\mathscr{D}_0$ is revision-proof.

Proof. Let $\widetilde{P}$ and a family of sets $\{\mathscr{D}_r\}$ satisfy the conditions in the proposition at some P^{t-1}. Suppose that σ is a revision-proof strategy and that $\Phi(\sigma|P^{t-1}) \in \mathscr{D}_0$. We obtain a contradiction. Let $\Phi(\sigma|P^{t-1}, \widetilde{P}^r) \in \mathscr{D}_r$. Since σ is revision-proof, it is also sub-game perfect (Proposition 1) and so:

$$U(t+r, P^{t-1}, \widetilde{P}^r, \Phi(\sigma|P^{t-1}, \widetilde{P}^r)) \geq U(t+r, P^{t-1}, \widetilde{P}^{r+1}, \Phi(\sigma|P^{t-1}, \widetilde{P}^{r+1})).$$

Thus, by condition (ii), $\Phi(\sigma|P^{t-1}, \widetilde{P}^{r+1}) \in \mathscr{D}_{r+1}$. Hence, by induction, for all $r \in \{0\} \cup \mathscr{N}$, $\Phi(\sigma|P^{t-1}, \widetilde{P}^r) \in \mathscr{D}_r$. Then, by condition (i), for all $r \in \{0\} \cup \mathscr{N}$,

$$U(t+r, P^{t-1}, \widetilde{P}) \geq U(t+r, P^{t-1}, \widetilde{P}^r, \Phi(\sigma|P^{t-1}, \widetilde{P}^r)).$$

with strict inequality at $r = 0$. By Lemma 1, this contradicts revision-proofness of σ. □

We give two applications of Proposition 3. Let $\mathscr{P}_t^*(P^{t-1}) := \text{argmax}_{\mathscr{P}^\infty} U(t, P^{t-1}, P)$ be the set of optimal plans for the t-th player in the sub-game $\mathscr{G}(U; P^{t-1})$. Let $\overline{U}_t(P^{t-1}) := \max_{P \in \mathscr{P}^\infty} U(t, P^{t-1}, P)$ be the associated optimal payoff.

Definition 5. U exhibits *weak agreement over optima* at t and P^{t-1}, if there is a P^* such that for all $r \in \mathscr{N}$,

$$P_r^* \in \mathscr{P}_{t-1+r}^*(P^{t-1}, P^{*r-1}).$$

U exhibits weak agreement over optima if it exhibits weak agreement over optima at all t and P^{t-1}.

If weak agreement over optima occurs at (t, P^{t-1}), then there is *a* P^* such that all players after (t, P^{t-1}) agree that it is better to continue with P^* than to switch to an alternative path. The agreement need concern only a single optima at (t, P^{t-1}). There may be other optimal paths for the player at t whose continuations are not optimal for successors. Alternatively, successor players may have other optimal continuation paths that are not optimal for the player at t. In this sense the agreement is weak. Weak agreement over optima is much weaker than time consistency which requires successive players to have identical preference orderings over continuation paths.

It is clear that weak agreement over optima is necessary for any strategy and, hence, necessary for all revision-proof strategies, to attain the optimal payoffs $\overline{U}_t(P^{t-1})$ in all sub-games.[11] Our first application of Proposition 3 shows that it is also *sufficient* for all revision-proof strategies to attain optima in all sub-games.

Corollary 1. *Assume that U exhibits weak agreement over optima. σ is revision-proof for $\mathscr{G}(U)$ if and only if for all $t \in \mathscr{N}$ and $P^{t-1} \in \mathscr{P}^{t-1}$, $\Phi(\sigma|P^{t-1})$ is in $\mathscr{P}_t^*(P^{t-1})$.*

Proof. If for all $t \in \mathscr{N}$ and $P^{t-1} \in \mathscr{P}^{t-1}$, $\Phi(\sigma|P^{t-1})$ is in $\mathscr{P}_t^*(P^{t-1})$, then evidently σ is revision-proof since then it is not possible to find any t, P^{t-1} and P' such that $U(t, P^{t-1}, P') > V_t(P^{t-1}, \sigma)$.

Conversely, suppose that σ is revision-proof and that there is some t and P^{t-1} such that $\Phi(\sigma|P^{t-1}) \notin \mathscr{P}_t^*(P^{t-1})$. Since U exhibits weak agreement over optima, there is a path $P^* \in \mathscr{P}_t^*(P^{t-1})$. Thus, $\Phi(\sigma|P^{t-1}) \in \mathscr{D}_0 := \{P|U(t, P^{t-1}, P^*) > U(t, P^{t-1}, P)\}$. Weak agreement over optima implies that for all $r \in \mathscr{N}$, $\mathscr{D}_r := \{P|U(t+r, P^{t-1}, P^*) \geq U(t+r, P^{t-1}, P^{*r}, P)\} = \mathscr{P}$. Thus, the conditions of Proposition 3 hold and σ is not revision-proof. □

Examples 1 and 2 show that weak agreement over optima is not sufficient to ensure that all sub-game perfect strategies attain optima in all sub-games. A sufficient condition for this is time consistency and continuity at infinity of U.

Definition 6. U is *time consistent* if for all $t \in \mathscr{N}$, (P^t, P) and $(P^t, P') \in \mathscr{P}^\infty$,

$$U(t, P^t, P) > U(t, P^t, P') \Leftrightarrow U(t+1, P^t, P) > U(t+1, P^t, P').$$

Definition 7. $U(t, P^{t-1}, \cdot)$ is *continuous at infinity* if it is continuous in the relative product topology on $\Pi_t(P^{t-1})$, i.e. if for each sequence $\{P(n)\}_{n=1}^\infty$ with $P(n) = \{p_r^n\}_{r=1}^\infty \in \Pi_t(P^{t-1})$, for all n, $p_r^n \to p_r$ and $P = \{p_r\}_{r=1}^\infty \in \Pi(P^{t-1})$, $\lim_{n\to\infty} U(t, P^{t-1}, P(n)) = U(t, P^{t-1}, P)$.

Proposition 4. *Let Assumption 1 hold. Assume that U is time consistent and that each $U(t, P^{t-1}, \cdot)$, $t \in \mathscr{N}$, is continuous at infinity, then a strategy σ is sub-game perfect for $\mathscr{G}(U)$ if and only if each $(\sigma|P^{t-1})$ attains the payoff $\overline{U}(t, P^{t-1})$.*

Proof. See Appendix. □

[11] If weak agreement did not hold, there would be at least one sub-game in which the attainment of an optimal payoff $\overline{U}_t(P^{t-1})$ by a player precludes its attainment by a later player.

Sub-game perfection imposes a "local" optimality requirement on dynasties (each player does the best she can given future play). When U exhibits time consistency and continuity at infinity, this local optimality translates into "global" optimality over continuation paths. Example 3 shows that time consistency is not sufficient to ensure that all sub-game perfect strategies attain optima in all sub-games. Since time consistency implies weak agreement over optima, the (unique) revision-proof strategy in the example does attain these optima.

Example 3. Let $\mathscr{P} = \mathbb{R}_+$, $p_0 \in \mathbb{R}_+$ be given and $U(t,P) = \liminf_{T\to\infty} \sum_{r=1}^{T} \beta^r [p_{r-1} - \delta p_r]$ with $\delta > \beta$. Clearly, for all t, $\overline{U}(t, P^{t-1}) = p_{t-1}$, but sub-game perfect strategies with $V_t(P^{t-1}, \sigma) < \overline{U}(t, P^{t-1})$ for all t and P^{t-1} are possible. These inflict ever increasing continuation penalties on players who fail to choose non-zero actions.

As a second application of Proposition 3, we give an example in which U does not exhibit weak agreement over optima and is not continuous at infinity. In the example, revision-proofness excludes paths whose sub-game perfect implementations require "explosively" bad sequences of penalties to deter finite length revisions

Example 4. Let $\mathscr{P} = \mathbb{R}$ and for all t and P, let:

$$U(t,P) = W(P_t) := (1-\beta)u(p_t) + \sum_{r=1}^{\infty} \beta^r (1-\beta) v(p_{t+r}),$$

where $P_t = \{p_{t-1+r}\}_{r=1}^{\infty}$. Assume that $u, v : \mathscr{P} \to \mathbb{R}$ are not equal, bounded above and unbounded below. Let $\widehat{p} \in \operatorname{argmax}_{\mathscr{P}} u(p)$ and $\underline{p} \in \operatorname{argmin}_{\mathscr{P}} v(p) - u(p)$, so that $\widehat{p}$ maximizes a player's current payoff and $\underline{p}$ minimizes the difference between a predecessor's continuation payoff v and a player's current payoff u. For each $r \in \{0\} \cup \mathscr{N}$, $\mathscr{D}_r = \mathscr{D} := \{P | u(\widehat{p}) + \beta[v(\underline{p}) - u(\underline{p})] > W(P)\}$. This set is non-empty since u and v are unbounded below. We show that *no* path in $\mathscr{D}$ is revision-proof.

Let $\widehat{P} = (\widehat{p}, \widehat{p}, \ldots)$. By simple algebra, $(1-\beta)u(\widehat{p}) + \beta v(\widehat{p}) > u(\widehat{p}) + \beta[v(\underline{p}) - u(\underline{p})]$ and so $\mathscr{D} \subset \{P | W(\widehat{P}) > W(P)\}$. Thus, the sequence of sets $\{\mathscr{D}_r\}$, each $\mathscr{D}_r = \mathscr{D}$, satisfies (i) in Proposition 3. It remains to check that (ii) in Proposition 3 holds. To this end, suppose $P \in \mathscr{D}$ and $W(P) \geq W(\widehat{p}, P')$. Thus, $(\widehat{p}, P') \in \mathscr{D}$ as well and so:

$$u(\widehat{p}) + \beta[v(\underline{p}) - u(\underline{p})] > (1-\beta)u(\widehat{p}) + \sum_{s=1}^{\infty} \beta^s (1-\beta) v(p'_s).$$

Rearranging the last expression and using $v(p'_1) - u(p'_1) \geq v(\underline{p}) - u(\underline{p})$ gives:

$$u(\widehat{p}) + \beta[v(\underline{p}) - u(\underline{p})] > (1-\beta)u(p'_1) + \sum_{s=1}^{\infty} \beta^s (1-\beta) v(p'_{s+1}),$$

i.e. $P' \in \mathscr{D}$. It then follows from Proposition 3 that for all histories, no continuation path in $\mathscr{D}$ is revision-proof. However, it is easy to check that $\mathscr{D}$ *does* contain paths that are sub-game perfect. These paths are sustained by reversion to ever more severe paths within $\mathscr{D}$. For example, if a player receives the payoff w with $u(\widehat{p}) + \beta[v(\underline{p}) - u(\underline{p})] - \varepsilon > w$, then defection from w can only be sustained by reversion to a path P' that delivers a payoff of w' to the subsequent player where:

$$\begin{aligned} u(\widehat{p}) + \beta[v(\underline{p}) - u(\underline{p})] - \varepsilon > w &\geq (1-\beta)u(\widehat{p}) + (1-\beta)\beta(v(p'_1) - u(p'_1)) + \beta w' \\ &\geq (1-\beta)\{u(\widehat{p}) + \beta(v(\underline{p}) - u(\underline{p}))\} + \beta w' \end{aligned}$$

so $w - \frac{(1-\beta)}{\beta}\varepsilon > w'$. Consequently, in a sub-game perfect equilibrium, an exploding sequence of punishments is necessary to deter long sequences of deviations from a path with initial payoff below $u(\widehat{p}) + \beta[v(\underline{p}) - u(\underline{p})]$. However, such paths cannot be implemented by a revision-proof strategy. $\square$

Suppose for any path in a set $\mathscr{D}_0$ it is possible to construct a strategy that (i) implements the path and (ii) given *any* candidate revision path, both deters successive player from being the last to join the revision and builds up their payoffs along the revision so that at least one obtains more from reversion to the strategy than from the complete revision. By definition that the paths in $\mathscr{D}_0$ will be revision-proof. We pursue this logic in the next section in the context of quasi-recursive games. We give conditions under which revision-proofness is relatively permissive with respect to paths. However, even in these cases, it places restrictions on "off-equilibrium" play and, hence, refines the set of sub-game perfect strategies.

5 Quasi-recursive Games

Many dynastic games have a quasi-recursive structure. See, inter alia, the diverse contributions of Bernheim, Ray, and Yeltekin (1999), Leininger (1986), Asheim (1997) and the examples below drawn from Kocherlakota (1996) and Hammond (1975). In these games a player and her successor have differing preferences over the successor's current action, but conditional on this action identical preferences over the successor's future action path. This section explores the implications of revision-proofness for play in quasi-recursive settings without state variables. In the next section, state variables are introduced.

For all $t \in \mathscr{N}$ and $P \in \mathscr{P}^{\infty}$, let

$$U(t, P) = W(P_t),$$

where $W : \mathscr{P}^{\infty} \to \mathbb{R}$ is constructed from a current payoff function $u : \mathscr{P} \to \mathbb{R}$ and a pair of intertemporal payoff aggregators $Q : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ and $R : \mathscr{P} \times \mathbb{R} \to \mathbb{R}$. Q gives a player's payoff by combining her current and continuation payoffs. R gives a player's continuation payoff by combining next period's action with next period's continuation payoff. Formally, suppose that $\mathscr{P}$ and the triple (u, Q, R) satisfy the following condition.

Assumption 2. *(i) $\mathscr{P}$ is a compact, convex subset of a normed space $(\mathscr{P}_0, \|\cdot\|)$. $\mathscr{P}^{\infty}$ is equipped with the associated relative product topology.*

(ii) $Q : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is increasing in both arguments, concave and and continuous. $u : \mathscr{P} \to \mathbb{R}$ is concave and continuous.

(iii) $R : \mathscr{P} \times \mathbb{R} \to \mathbb{R}$ is non-decreasing in its second argument, concave and continuous on $\mathscr{P} \times \mathbb{R}$. Also, there is a $\beta \in [0, 1)$ such that for all $y, y' \in \mathbb{R}$,

$$\sup_{\mathscr{P}} |R(p, y) - R(p, y')| \le \beta |y - y'|$$

Under Assumption 2, it is straightforward to show that there is a unique bounded, continuous and concave function $Y : \mathscr{P}^{\infty} \to \mathbb{R}$ satisfying for all (p, P), $Y(p, P) = R(p, Y(P))$. We refer to Y as a *continuation* payoff function and define the payoff function $W : \mathscr{P}^{\infty} \to \mathbb{R}$ according to for all (p, P), $W(p, P) = Q(u(p), Y(P))$. Under Assumption 2, W is also bounded, continuous and concave. This specification implies that given a history P^{t+r} the t-th and $t+r$-th players have the same preferences over continuation paths P^{t+r+1}. In contrast the stronger assumption of time consistency requires that given a history P^{t+r-1}, the t-th and $t+r$-th players have the same preferences over continuation paths P^{t+r}.

The quasi-recursive structure accommodates many types of time inconsistency considered in the literature. We give two concrete applications below. Subsequently, we use the first to illustrate the basic logic of our results. The second shows how a policy game may be mapped into the quasi-recursive framework.

Example *(Hammond's overlapping generations pension game, see Hammond (1975)).* A two period-lived agent is born at each $t \in \mathcal{N}$; a one period-lived old agent lives at $t = 1$. Agents have an endowment of 1 when young and 0 when old. A young agent may share her endowment with a contemporaneous old agent. Let $u^y, u^o : \mathbb{R}_+ \to \mathbb{R}$ denote the utility from consumption of young and old agents. Players are identified with two period lived agents and are indexed by the date $t \in \mathcal{N}$ at which they are born and are young. Let p_t denote the consumption of the player (young agent) at t. The t-th player's preferences over paths (of young agent consumption) can be expressed as $U(t, P) = u^y(p_t) + u^o(1 - p_{t+1})$. Thus, $u = u^y$, $Q(u(p), y) = u(p) + y$ and $R(p, y) = u^o(1 - p)$. $\square$

Example *(Kocherlakota's policy game, see Kocherlakota (1996)).* A two period-lived household is born at each $t \in \mathcal{N}$; a one period-lived old household lives at $t = 1$. Households consume a quantity of a private good c when young and a quantity of a public good g when old. The household born in period t has preferences: $u_t(c_t, g_{t+1}) = c_t + \frac{1}{2} g_{t+1}$ over private and public consumption in periods t and $t + 1$. Households receive an endowment of income q when old and must borrow against this to finance private consumption when young. They can borrow on an international loan market at rate r. At each date $t \in \mathcal{N}$, a government taxes the income of the old at rate p_t to finance public consumption. The young household at $t \in \mathcal{N}$ borrows the maximal amount subject to being able to pay her taxes, i.e. borrows $(1 - p_{t+1})q/(1 + r)$. The goal of a date t government is to maximize the discounted utility of current and future generations of households including the current old generation, i.e. to maximize:

$$\sum_{r=0}^{\infty} \beta^r \left(c_{t+r} + \frac{1}{2} g_{t+r} \right). \tag{5}$$

By substituting optimal household choices and the government budget constraints $g_{t+r} = p_{t+r} q$ into (5), the objective of the government at t can be re-expressed as a function of the continuation tax path P_t:

$$W(P_t) = \sum_{r=0}^{\infty} \beta^r q \left(\frac{1}{2} p_{r+t} + \frac{(1 - p_{t+r+1})}{(1 + r)} \right). \tag{6}$$

Under the assumption that $r < 1$, the preferred path for the government at t is $p_t = 1$, because she derives no utility from the past consumption of the present old, and $p_{t+r} = 0$, $r \in \mathcal{N}$, to enable high private consumption of current and future young. Rearranging (6) gives:

$$W(P_t) = \frac{1}{2} p_t q + \frac{q}{1 + r} \frac{1}{1 - \beta} - \left(\frac{1}{1 + r} - \frac{\beta}{2} \right) q \sum_{r=1}^{\infty} \beta^{r-1} p_{t+r}.$$

By setting $u(p) = \frac{1}{2}pq + \frac{q}{1+r}\frac{1}{1-\beta}$, $Q(u(p),y) = u(p) + \beta y$, $v(p) = -\left(\frac{1}{1+r} - \frac{\beta}{2}\right) qp$ and $R(p,y) = v(p) + \beta y$, a quasi-recursive game may be associated with the example. □

5.1 Sub-game perfection in quasi-recursive games

As a precursor to subsequent analysis of revision-proofness, we provide some characterization of sub-game perfect strategies in quasi-recursive games. Let $\mathscr{Y}$ denote the set of sub-game perfect continuation payoffs,

$$\mathscr{Y} := \{y | y = Y(\Phi(\sigma)), \quad \sigma \text{ is sub-game perfect}\}.$$

Proposition 11 in Appendix B shows that under Assumption 2, $\mathscr{Y}$ is a compact interval, $[\underline{y}, \overline{y}]$. Lemma 3 in the same appendix characterizes the endpoints of this interval. It shows that there is a stationary sub-game perfect action path $(\overline{p}, \overline{p}, \ldots)$ that attains $\overline{y}$ and satisfies $\overline{y} = R(\overline{p}, \overline{y})$. It shows that for the worst sub-game perfect continuation payoff $\underline{y}$ there are two possibilities. In the first, $\underline{y}$ is obtained by a stationary sub-game perfect action path $(\underline{p}, \underline{p}, \ldots)$, $\underline{y} = R(\underline{p}, \underline{y})$ and $u(\underline{p}) = u^* := \max_{\mathscr{P}} u(p)$. In this case, action $\underline{p}$ is best for a current player and, amongst actions consistent with sub-game perfection, worst for her predecessor. For the second case, there is no stationary sub-game perfect path that attains the worst sub-game perfect continuation payoff. Instead this payoff is attained by path that begins with a severe action $\underline{p}$ satisfying $\underline{y} > Y(\underline{p}, \underline{p}, \ldots)$. The subsequent period's continuation payoff $y'(\underline{p})$ satisfies $\underline{y} = R(\underline{p}, y'(\underline{p})) < y'(\underline{p})$.

The previous remarks relate to sub-game perfect *continuation payoffs*. The best and worst sub-game perfect *payoffs* for a player are given by $Q(u^*, \overline{y})$ and $Q(u^*, \underline{y})$.

Example *(Hammond's pension game, cont.)* The action $p_t = 1$ is both the best current action for the t-th player, since it implies she consumes the entire endowment when young, and the worst continuation action for the $t-1$-th player since it implies she gets nothing when old. Consequently, the "no-sharing path" $\underline{P} = (1, 1, 1, \ldots)$ gives both the worst sub-game perfect payoff $Q(u^*, \underline{y})$ and the worst sub-game perfect continuation payoff $\underline{y}$ to all players. The best sub-game perfect continuation payoff $\overline{y}$ is attained by a path on which a player's successor consumes the minimal amount consistent with sub-game perfection, i.e. consumes $\overline{p}$, where:

$$\overline{p} = \min_{(p,p') \in \mathscr{P}^2} \{p | u^y(p) + u^o(1-p') \geq u^y(1) + u^o(0), p' \in [\overline{p}, 1]\}. \tag{7}$$

There are two possibilities. In the first, $\overline{p} = 0$ and $u^y(0) + u^o(1) \geq u^y(1) + u^o(0)$, i.e. the player's successor consumes nothing when young and gives everything to her predecessor. She in turn receives a large amount (possibly the entire endowment) when old. In the second case, $\overline{p} \in (0, 1]$ and $u^y(\overline{p}) + u^o(1 - \overline{p}) = u^y(1) + u^o(0)$. In this case, the player's successor gives the maximal amount $1 - \underline{p}$ consistent with sub-game perfection when young and is motivated to do so by a gift of this amount when old. $(1, \overline{p}, \overline{p}, \ldots)$ is the unique best sub-game perfect path for a player. $\square$

5.2 Revision-proofness in quasi-recursive games

We now consider revision-proofness in quasi-recursive games. To build intuition consider a finite horizon game lasting T periods. Suppose that the T-th player has an objective $W_T : \mathscr{P} \to \mathbb{R}$ and the $T-1$-th player a continuation objective $Y_T : \mathscr{P} \to \mathbb{R}$. Further suppose that prior players, $t = 1, \ldots, T-1$, have objectives and continuation objectives $W_t : \mathscr{P}^{T+1-t} \to \mathbb{R}$ and $Y_t : \mathscr{P}^{T+1-t} \to \mathbb{R}$ generated by increasing aggregators Q and R and current payoff functions $u, v : \mathscr{P} \to \mathbb{R}$ according to $W_t(P^{T+1-t}) = Q(u(p_t), Y_{t+1}(P^{T-t}))$ and $Y_t(P^{T+1-t}) = R(v(p_t), Y_{t+1}(P^{T-t}))$.[12] In this setting a revision-proof strategy σ can be derived by a backwards induction argument. Specifically, for all histories P^{T-1}, let σ_T be such that:

$$\sigma_T(P^{T-1}) \in \underset{\mathscr{U}_T^*}{\operatorname{argmax}}\, Y_T(p), \quad \mathscr{U}_T^* = \underset{\mathscr{P}}{\operatorname{argmax}}\, W_T(p).$$

Thus, σ_T maximizes the $T-1$-th player's continuation payoff Y_T subject to maximizing the T-th player's payoff W_T. For prior dates t and histories P^{t-1}, let σ_t be such that:

$$\sigma_t(P^{t-1}) \in \underset{\mathscr{U}^*}{\operatorname{argmax}}\, v(p), \quad \mathscr{U}^* = \underset{\mathscr{P}}{\operatorname{argmax}}\, u(p). \tag{8}$$

Then each $\sigma_t(P^{t-1})$ maximizes the continuation payoff v of the $t-1$-th player subject to maximizing the current payoff u of the t-th player. A strategy constructed in this way prescribes optimal actions and breaks indifferences over a current player's optimal actions in favor of prior players. Consequently, it is not possible to raise the payoff of an earlier player without reducing the payoff of a later one. The strategy is revision-proof.

The above construction relies on a terminal period in which the last player obtains her optimal payoff. The argument cannot be applied to infinite horizon games, but it

[12] Here, in a slight abuse of notation, we simplify by assuming that R aggregates a current payoff $v(p_t)$ rather than a current action p_t with a continuation payoff. Hence, R is fully separable.

suggests a procedure that can. Recall that in finite horizon games, a revision-proof strategy prescribes an action for the terminal player that is best for her and, subject to this, best for her predecessors. Rather trivially, this player "blocks" attempts by predecessors to revise play in the terminal period. For infinite horizon games, we design revision-proof strategies that endogenously create "blocking players" in all sub-games. These players are analogous to terminal players in finite horizon games. They receive a best sub-game perfect payoff and take actions that are maximal for predecessors subject to being maximal for themselves. The payoffs of a blocking player or her successors are reduced by any modification to their play that raises predecessor payoffs. Thus, blocking players create obstacles to revision. We use Hammond's pension game to show how revision-proof strategies that feature blocking players can be constructed. In this particular game, *all* sub-game perfect paths can be implemented with such strategies and so all are revision-proof.

Example *(Hammond's pension game, cont.)* Let $\widetilde{P}$ be an arbitrary sub-game perfect path. Consider the strategy σ defined as follows. First, $\Phi(\sigma) = \widetilde{P}$. Second, for all t and $P^t \neq (P^{t-1}, \sigma_t(P^{t-1}))$, $\Phi(\sigma|P^t) = (\widehat{p}, \overline{p}, \overline{p}, \ldots)$, where $\widehat{p} = 1$ and $\overline{p}$ is defined in (7). Thus, the player following a defection takes her unique best current action and obtains her best sub-game perfect payoff. In the language of the previous paragraph, this player is a blocking player.

It is easy to verify that σ is sub-game perfect; we check that it is revision-proof. We focus on the simplest case in which $\overline{p} = 0$, i.e. the case in which it is possible to sustain 0 consumption when young as part of a sub-game perfect strategy, because players sufficiently value consumption when old.[13] Let P^t be a history and P'' a successful revision path at P^t. Then, each player $t + r$, $r \in \mathcal{N}$, receives a weakly higher payoff and some receive a strictly higher payoff from P'' than from reversion to the strategy. Thus, for all $r \in \mathcal{N}$, $u^y(p''_r) + u^o(1 - p''_{r+1}) \geq W(\Phi(\sigma|P^t, P''^{r-1}))$ with the inequality strict for some r. Let $t + r_0$ be the first date such that $u^y(p''_{r_0}) + u^o(1 - p''_{r_0+1}) > W(\Phi(\sigma|P^t, P''^{r_0-1}))$. Then either p''_{r_0} is above or p''_{r_0+1} is below the corresponding actions prescribed by σ. Suppose that $p''_{r_0} \neq \sigma_{t+r_0}(P^t, P''^{r_0-1})$. Then $\Phi(\sigma|P^t, P''^{r_0}) = (1, 0, 0, \ldots)$. This continuation path is not just the best sub-game perfect path for the $t + r_0 + 1$-th player, it is a best possible path for this player: she consumes the maximal amount of 1 in both periods of her life. The player at $t + r_0 + 1$ will "block" revisions that fail to match this. Thus, if $p''_{r_0} \neq \sigma_{t+r_0}(P^t, P''^{r_0-1})$, then p''_{r_0+1} must equal 1 and p''_{r_0+2} must equal 0. Hence, the $t + r_0$-

[13] Our analysis extends to situations in which $\overline{p} > 0$ and this case is covered by our general result, Proposition 5. However, to maximize transparency, we do not detail it here.

th player receives a payoff of $u^y(p''_{r_0}) + u^o(1 - p''_{r_0+1}) = u^y(p''_{r_0}) + u^o(0) \leq u^y(1) + u^o(0)$. But $u^y(p''_{r_0}) + u^o(1 - p''_{r_0+1}) > W(\Phi(\sigma|P^t, P''^{r_0-1})) \geq u^y(1) + u^o(0)$, where the first inequality is by assumption and the second stems from the fact that σ is sub-game perfect and $u^y(1) + u^o(0)$ is the lowest possible sub-game perfect payoff. However, this leads to a contradiction. So, p''_{r_0} must equal and p''_{r_0+1} must be lower than the actions prescribed by the strategy. In particular, $p''_{r_0+1} < 1$. But exactly the same argument applied at $t + r_0 + 1$ shows that $p''_{r_0+1} = 1 \geq \sigma_{t+r_0+1}(P^t, P''^{r_0})$, another contradiction. We conclude that there are no successful revision paths at any history. Thus, by Lemma 1, σ and, hence, $\widetilde{P}$ is revision-proof. Since $\widetilde{P}$ was an arbitrary sub-game perfect path, it follows that all such paths are revision-proof in this game. $\square$

The pension game has two features that greatly simplify the analysis. First, players care only about actions in the two periods of their life and, second, a player and her immediate successor have strictly conflicting preferences over the successor's action. Consequently, the unique optimal action for a successor gives the worst possible continuation payoff for a predecessor. This enables a blocking player to be placed in every sub-game immediately following a deviation. Proposition 5 shows that the essential logic of the example generalizes. In more general settings, a blocking player who attains her best sub-game perfect payoff may not inflict a very severe penalty on her immediate predecessors. Hence, to sustain a given path, it is often necessary to embed the blocking player deeper into the sub-game following a defection. We construct (revision-proof) strategies that do this. These strategies prescribe repeated play of an action $\widehat{p}$ that maximizes a player's current payoff function u and, given this, maximizes the continuation payoff of predecessors. Thus, $\widehat{p}$ solves a problem similar to (8) in the finite horizon setting. By making the string of $\widehat{p}$ plays long enough a continuation payoff very nearly equal to $\widehat{y} = R(\widehat{p}, \widehat{y})$, $\widehat{y} = Y(\widehat{P})$, $\widehat{P} = (\widehat{p}, \widehat{p}, \ldots)$ is attained. This is more severe than $Y(\widehat{p}, \overline{P}) = R(\widehat{p}, \overline{y})$, the continuation payoff obtained if a blocking player immediately follows a defector, and, hence, sustains more paths. As time passes and the blocking player is approached player continuation payoffs are built up towards $\overline{y}$. The blocking player or her successors are made worse off by any revision to their play that benefits predecessors. Given this, the player immediately prior to the blocking player cannot revise future play in her favor and by playing $\widehat{p}$ obtains her best possible current payoff. Further any change to her play that benefits prior players makes her worse off. This logic extends back through the game tree until the initial defection that triggered play of the $\widehat{p}$ actions is reached. If the initial defecting player is better off adhering to the strategy than triggering this play, then the strategy is proof against any revision that benefits her and, conditional on her defection, players between her and the blocking player. This leaves open the possibility of a successful revision for players who move after the blocking player. Such a possibility is eliminated if the continuation of a best sub-game perfect path can itself be sustained by reversion to play involving a later blocking player.

Remark 1. Analysis of the pension game was simplified by the fact that a player and her successor had conflicting objectives over the later player's action and that players cared about actions in only two periods. Kocherlakota (1996)'s policy game has the first of these properties, but not the second. Pursuing the logic sketched above and formalized in Proposition 5 below, as in the pension game, all sub-game perfect paths may be shown to be revision-proof in this game.[14]

We impose two assumptions. The first is on the aggregator R and uses the following definition.

Definition 8. Define the set of best current actions by $\mathscr{U}^* := \text{argmax}_{\mathscr{P}} u(p)$ and the correspondence of best current actions that are best for predecessors by $\widetilde{\mathscr{R}}^* : Y(\mathscr{P}^\infty) \to 2^{\mathscr{P}}$, $\widetilde{\mathscr{R}}^*(y) := \text{argmax}_{\mathscr{U}^*} R(p, y)$.

Assumption 2 ensures that $\mathscr{U}^*$ and each $\widetilde{\mathscr{R}}^*(y)$ is non-empty and compact.

Assumption 3. *The correspondence $\widetilde{\mathscr{R}}^*$ has a constant value $\mathscr{R}^*$.*

Assumption 3 holds if the restriction of R to $\mathscr{U}^* \times \mathscr{Y}$ has the form $R(p, y) = R'(v(p), y)$ with $v : \mathscr{U}^* \to \mathbb{R}$ and each $R'(\cdot, y)$ increasing on $v(\mathscr{U}^*)$. This in turn holds trivially if $\mathscr{U}^*$ is single-valued or if R has the form $R'(v(p), y)$ on all of $\mathscr{P} \times \mathscr{Y}$. Let $\widehat{p}$ belong to $\mathscr{R}^*$ and define $\widehat{P} = (\widehat{p}, \widehat{p}, \ldots)$, $\widehat{w} := Q(u(\widehat{p}), Y(\widehat{P}))$ and $\widehat{y} := Y(\widehat{P})$. In quasi-recursive games with weak agreement over optima $\widehat{w}$ is the best sub-game perfect payoff for a player; in Hammond's pension game, Kocherlakota's policy game and games in which all best current actions are worst for predecessors $\widehat{w}$ is the worst sub-game perfect payoff.

We now provide our second assumption.

Assumption 4. *$Q(u(\overline{p}), \overline{y}) \geq \widehat{w} := Q(u(\widehat{p}), \widehat{y})$. If $Q(u(\overline{p}), \overline{y}) = \widehat{w} := Q(u(\widehat{p}), \widehat{y})$, then $R(u(\overline{p}), \overline{y}) \geq \widehat{y} := R(u(\widehat{p}), \widehat{y})$.*

This assumption implies that a player weakly prefers implementing the continuation of a predecessor's preferred sub-game path $\overline{P}$ (and obtaining her own best continuation

[14] More precisely, Proposition 5 applies to the reduced form quasi-recursive game that we associated with Kocherlakota's policy game. The former is reduced form in that the (optimal) actions of private households are substituted out. The revision-proof refinement applied to this game translates into a revision-proof refinement for the fully specified game played by policymakers *and* private households.

payoff $\overline{y}$) to taking a best current action that is best for predecessors $\widehat{p}$ and having all successors do the same. Further, if a player is indifferent between these paths, then her predecessor weakly prefers the former. We use this assumption to ensure that neither players nor, if they are indifferent, their predecessors prefer deviating from the best sub-game perfect path used to reward a blocking player to paths with payoffs approximately equal to $\widehat{w}$ that feature a later blocking player. Assumption 4 holds automatically in Hammond's pension game and Kocherakota's policy game since in these games $\widehat{P}$ is a worst sub-game perfect path and, hence, gives a payoff $\widehat{w}$ below that from the sub-game perfect path $\overline{P}$. It also holds in games with weak agreement over optima in which paths of the form $\widehat{P}$ are simultaneously best and best continuation sub-game perfect paths.

We now state our first main result for quasi-recursive games.

Proposition 5. *Let Assumptions 2 to 4 hold. Let $\widetilde{P}$ be any path such that (i) for all t, $W(\widetilde{P}_t) \geq \widehat{w}$ and (ii) if $t > 1$ and $W(\widetilde{P}_t) = \widehat{w}$, then $Y(\widetilde{P}_t) \geq \widehat{y}$. There is a revision-proof strategy σ that implements the path $\widetilde{P}$.*

Proof. See Appendix C. □

Corollary 2 below strengthens condition (ii) of Proposition 5 to hold in all periods including the first and obtains a stronger result. It is now not possible to raise the *first period* continuation payoff $Y(\widetilde{P})$ implied by the strategy without reducing the payoff of a later player. This stronger conclusion is not needed to establish revision-proofness since there is no 0-th period player to receive a continuation payoff of $Y(\widetilde{P})$. However, it is useful in deriving later results.[15]

Corollary 2. *Let Assumptions 2 to 4 hold. Let $\widetilde{P}$ be any path such that (i) for all t, $W(\widetilde{P}_t) \geq \widehat{w}$ and (ii) for all t, if $W(\widetilde{P}_t) = \widehat{w}$, then $Y(\widetilde{P}_t) \geq \widehat{y}$. There is a revision-proof strategy σ that implements the path $\widetilde{P}$ and there is no alternative σ' that delivers the same payoffs to all players as σ and satisfies $Y(\Phi(\sigma')) > Y(\widetilde{P})$.*

Proof. See Appendix C. □

In Hammond's pension game and Kocherlakota's policy game the worst sub-game perfect path involves repeated play of the unique best current action $\widehat{p} = 1$. Proposition 5 implies that in such games *all* sub-game perfect paths are revision proof. However, it is easy to construct quasi-recursive games in which repeated play of a best current action is *not* a worst sub-game perfect path. In such games, there are sub-game perfect paths with

[15]Specifically, it is used in the proof of Theorem 2, where a revision-proof strategy is constructed from revision-proof continuation strategies.

payoffs below $\widehat{w}$ at some or all dates. Proposition 5 is silent on the revision-proofness of these.

We now show that with some strengthening of the conditions of Proposition 5, *all* paths delivering payoffs strictly above the worst sub-game perfect payoff, $\underline{w}$, are revision-proof.[16] To the extent that these conditions hold, the revision-proof concept is quite permissive with respect to paths and outcomes. The argument underpinning the result is quite long, but the idea is simple. Recall that $\underline{p}$ is the first action of a sub-game perfect path giving the lowest continuation payoff $\underline{y}$, see Subsection 5.1. Given a path $\widetilde{P}$ that delivers payoffs $W(\widetilde{P}_t) > \underline{w}$ to successive players, a strategy σ is constructed that implements this path and prescribes $\underline{p}$ following a defection (from $\widetilde{P}$) until player payoffs are driven above $\widehat{w}$. Irrespective of post-defection play $\widehat{w}$ is achieved within a finite number of periods. Once $\widehat{w}$ is reached a revision-proof continuation strategy defined as in Corollary 2 is pursued. Figure 1 illustrates the construction.

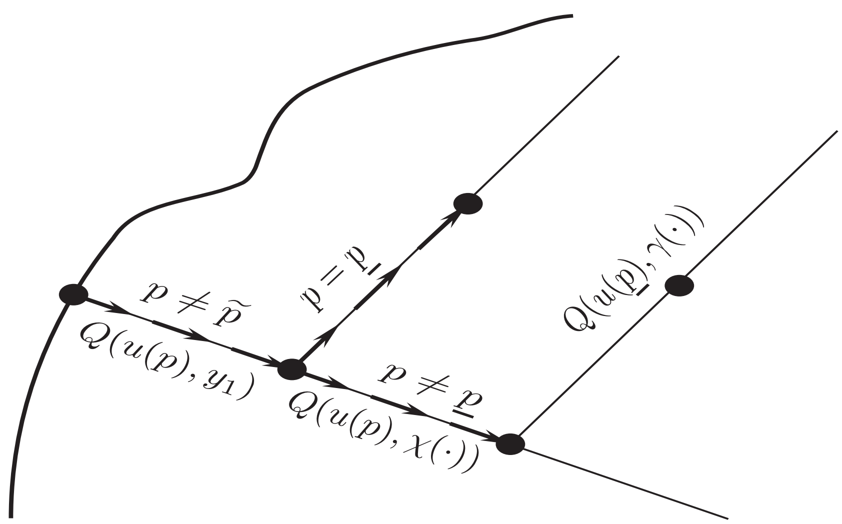

Figure 1: Construction of revision-proof σ. Following play of $p \neq \widetilde{p}$, the continuation strategy prescribes play of a sequence of $\underline{p}$ actions until a sub-game is reached in which a revision-proof continuation strategy is played. This play implies a continuation payoff of y_1 to the initial defecting player, sufficiently low to deter unilateral defection. Along the $\underline{p}$ path, continuation payoffs evolve according $y_{t+1} = \gamma(y_t) > y_t$. If a further defection occurs to $p \neq \underline{p}$, reversion to play of $\underline{p}$ and, eventually, a continuation revision-proof strategy is prescribed. This play delivers a continuation payoff $\chi(y_1) > y_1$ to the defecting player that strictly deters unilateral defection. Following plays of $\underline{p}$ or defections, continuation payoffs and, hence, payoffs are built up until the latter exceed $\widehat{w}$ and a continuation revision-proof strategy is played.

[16]If $\underline{w} < \widehat{w}$, then paths with payoffs that remain above $\underline{w}$, but are sometimes below $\widehat{w}$ are revision-proof.

The construction of σ implies sub-games in which revision-proof continuation strategies are played. By definition, it is not possible to revise these so as to raise the payoffs of a player in the sub-game without reducing the payoff of some other player in the sub-game. Revisions that begin and end before such a revision-proof continuation strategy is reached (and player payoffs have been driven above $\widehat{w}$) involve a last deviator whose payoff is reduced by construction. Revisions that begin before a revision-proof continuation strategy is reached, but continue on into a sub-game in which such a continuation strategy is played are more complicated, but these too turn out to leave all players' payoffs unaltered or they reduce the payoffs to some.

Our argument requires two further assumptions. They are used to ensure that $\underline{p}$ prescriptions are consistent with the building up of payoffs to values above $\widehat{w}$. Assumption 5 asserts that a player is better off playing $\underline{p}$ and receiving a best sub-game perfect continuation payoff $\overline{y}$, than taking a best current action $\widehat{p}$ that is best for predecessors and having all her successors do the same. It combines concern with the future with disagreement between successive players about the later player's play (i.e. play of $\underline{p}$ by a successor is not too bad for the successor, but is bad for a predecessor; play of $\widehat{p}$ is optimal for a successor, but bad for a predecessor).

Assumption 5. $Q(u(\underline{p}), \overline{y}) > Q(u(\widehat{p}), \widehat{y}) = \widehat{w}$.

Assumption 6 restricts the slope of Q in its second argument.

Assumption 6. *For all $y \geq y'$ and p, there is a value $\kappa > 0$ such that*

$$\beta\kappa(y - y') \leq Q(u(p), y) - Q(u(p), y') < \kappa(y - y').$$

This assumption is satisfied if, in addition to Assumption 2, Q and R are quasi-linear. For example, if $Q(u(p), y) = u(p) + \delta y$ and $R(p, y) = v(p) + \beta y$ with $\delta \in (0, \infty)$ and $\beta \in (0, 1)$, then $Q(u(p), y) - Q(u(p), y') = \delta(y - y')$ and Assumption 6 holds with $\kappa = \delta/\beta$. We now state the second main result of this section.

Theorem 2. *Suppose that Assumptions 2 to 6 hold. Let $\widetilde{P}$ be any path such that for all t, $W(\widetilde{P}_t) > \underline{w}$, then P can be supported by a revision-proof strategy.*

Proof. In Appendix D. □

We give a simple example of a game that satisfies these assumptions.

Example *(Overlapping generations game with partial agreement)* Overlapping generations of players $t \in \mathcal{N}$ live for two periods. Let $\mathscr{P} = [0, 1]$. The objective of the t-th player, who is young in period t and old in $t+1$ is given by:

$$U(t, P) = u^y(p_t) + u^o(p_{t+1}),$$

where $u^y(p) = \frac{1}{2}p(1-p)$ and $u^o(p) = 1-p$. In terms of our general quasi-recursive notation: $u = u^y$, $Q(u(p), y) = u(p) + y$ and $R(p, y) = u^o(p)$.

For this example, u^y is maximized by $\widehat{p} = \frac{1}{2}$, while u^o is minimized by $\underline{p} = 1$ and maximized by $\overline{p} = 0$. $(\widehat{p}, \underline{p}, \overline{p}, \overline{p}, \ldots) = (1/2, 1, 0, 0, \ldots)$ is a worst sub-game perfect path. Its continuation $(1, 0, 0, \ldots)$ has a worst sub-game perfect continuation payoff $\underline{y} = 0$ and it has a payoff of $Q(u(\widehat{p}), \underline{y}) = \frac{1}{8}$. $(\widehat{p}, \overline{p}, \overline{p}, \overline{p}, \ldots) = (1/2, 0, 0, 0, \ldots)$ is a best sub-game perfect path. Its continuation $(0, 0, 0, \ldots)$ has a best sub-game perfect continuation payoff $\overline{y} = 1$ and it has a payoff of $Q(u(\widehat{p}), \overline{y}) = \frac{9}{8}$. Repetition of the unique best current action $\frac{1}{2}$ gives a payoff of $Q(u(\widehat{p}), \widehat{y}) = \frac{5}{8}$ which exceeds the worst sub-game perfect payoff. Assumption 4 is satisfied, $1 = Q(u(\overline{p}), \overline{y}) > Q(u(\widehat{p}), \widehat{y}) = \frac{5}{8}$. Thus, Proposition 5 applies and sub-game perfect paths with payoffs that remain above $\frac{5}{8}$ can be implemented by a revision-proof strategy. However, Proposition 5 does not apply to those sub-game perfect paths that deliver payoffs between $\frac{1}{8}$ and $\frac{5}{8}$ to some players. Assumption 5 is satisfied since $1 = Q(u(\underline{p}), \overline{y}) > \widehat{w} = Q(u(\widehat{p}), \widehat{y}) = \frac{5}{8}$. and the quasi-linearity of Q and R ensure that Assumption 6 holds. By Theorem 2 all paths with payoffs that remain strictly above $\frac{1}{8}$ are revision-proof. $\square$

6 Quasi-recursive Games With State Variables

We briefly extend the quasi-recursive environment of Section 5 to allow for state variables and, hence, history dependence in player preferences or constraints. Let $\mathscr{K} \subseteq \mathbb{R}^m$ denote a set of feasible states. Define a law of motion for states $\Lambda : \mathbb{R}^m \times \mathscr{P} \to \mathbb{R}^m$ and for $k \in \mathscr{K}$, let $\Gamma(k) = \{p | \Lambda(k, p) \in \mathscr{K}\}$. Assume that Γ is non-empty valued for all $k \in \mathscr{K}$. Let $\Lambda^1 = \Lambda$ and, for $t \geq 1$, let $\Lambda^{t+1} : \mathscr{K} \times \mathscr{P}^{t+1} \to \mathbb{R}^m$ be defined recursively according to:

$$\Lambda^{t+1}(k, P^t, p_{t+1}) = \Lambda(\Lambda^t(k, P^t), p_{t+1}). \tag{9}$$

$\Lambda^{t+1}(k, P^{t+1})$ is the period $t+2$ state given an initial state $k \in \mathscr{K}$ and choices P^{t+1} between periods 1 and $t+1$. The set of feasible paths given initial state $k \in \mathscr{K}$ is:

$$\Pi(k) = \{P \in \mathscr{P}^\infty | \forall t, \Lambda^t(k, P^t) \in \mathscr{K}\}. \tag{10}$$

Given an initial k, the function $U : \mathscr{N} \times \mathscr{P}^\infty \to \mathbb{R} \cup \{-\infty\}$ is generated according to:

$$U(t, P) = W(\Lambda^t(k, P^{t-1}), P_t), \tag{11}$$

where $W : \mathscr{K} \times \mathscr{P}^\infty \to \overline{\mathbb{R}}$, $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty\}$ is defined below. It follows from (11) that the t-th player's objective depends on histories P^{t-1} through their effect on the current state variable $k_t = \Lambda^t(k, P^{t-1})$. This formulation accommodates psychological state variables such as habit formation or altruism for predecessors as well as physical state variables via the assigning of $-\infty$ payoffs to infeasible choices. W is constructed from Λ, $u : \mathscr{K} \times \mathscr{P} \to \overline{\mathbb{R}}$ and aggregators $Q : \overline{\mathbb{R}} \times \overline{\mathbb{R}} \to \overline{\mathbb{R}}$ and $R : \mathscr{K} \times \mathscr{P} \times \overline{\mathbb{R}} \to \overline{\mathbb{R}}$. Similar to Section 5, Λ and R are used to construct a continuation payoff function $Y : \mathscr{K} \times \mathscr{P}^\infty \to \mathbb{R} \cup \{-\infty\}$. This is done formally in Theorem 3, Appendix E. Payoffs are constructed from Q, u and Y according to:[17]

$$W(k, p, P') = Q(u(k, p), Y(\Lambda(k, p), P')).$$

Let $\overline{\mu}(k) = \sup_{\Pi(k)} Y(k, P)$ and $\underline{\mu}(k) = \inf_{\Pi(k)} Y(k, P)$. It is easily seen that:

$$\overline{\mu}(k) = \sup_{p \in \Gamma(k)} R(k, p, \overline{\mu}(\Lambda(k, p))), \qquad \underline{\mu}(k) = \inf_{p \in \Gamma(k)} R(k, p, \underline{\mu}(\Lambda(k, p))).$$

Define the sub-game perfect continuation payoff correspondence $\mathscr{Y} : \mathscr{K} \to 2^{\mathbb{R}} \backslash \emptyset$,

$$\mathscr{Y}(k) = \{y | y = Y(k, \Phi(\sigma)), \sigma \text{ is sub-game perfect}\} \subseteq \mathscr{Y}_0(k) = [\underline{\mu}(k), \overline{\mu}(k)].$$

Proposition 12 in Appendix E asserts that $\mathscr{Y}$ is compact valued. Given this we can identify maximal and minimal sub-game perfect continuation payoff functions: $\overline{y}(k) = \max \mathscr{Y}(k)$ and $\underline{y}(k) = \min \mathscr{Y}(k)$. We can further show that:

$$\begin{aligned} \overline{y}(k) = \max R(k, p, \overline{y}(\Lambda(k, p))) \text{ s.t. } & p \in \Gamma(k), \\ & Q(u(k, p), \overline{y}(\Lambda(k, p))) \geq \sup_{p' \in \Gamma(k)} Q(u(k, p'), \underline{y}(\Lambda(k, p'))) \end{aligned} \tag{12}$$

and that the best possible sub-game perfect payoff at each $k \in \mathscr{K}$ is given by:

$$\overline{w}(k) = \sup_{\Gamma(k)} Q(u(k, p), \overline{y}(\Lambda(k, p))) \tag{13}$$

[17] It is easy to check that these definitions are consistent with Assumption 1 with $\Pi_t(P^{t-1}) = \Pi(\Lambda^t(k, P^{t-1}))$.

Let $\overline{p} : \mathscr{K} \to \mathscr{P}$ denote a selection from the optimal policy correspondence for (12). Let:

$$S(y)(k) = \max_{\mathscr{U}^*(y)(k)} R(k, p, y(\Lambda(k, p))),$$

where $\mathscr{U}^*(y)(k) = \text{argmax}_{\Gamma(k)} Q(u(k,p), y(\Lambda(k,p)))$. $S(y)(k)$ is the best possible continuation payoff for a preceding player at k given that the current player makes an optimal choice and is faced with the continuation payoff function y. Let $y_0 = \overline{y}$ and consider the iteration: $y_{n+1} = S(y_n)$, $n \in \{0\} \cup \mathscr{N}$. The continuation value function y_n is induced by a sequence of n players taking a best action for predecessors amongst their best action sets and the n-th player obtaining her best possible payoff. Thus, for some initial $k_0 \in \mathscr{K}$ and $T \in \mathscr{N}$, the policy correspondences associated with the maximizations $\max_{\mathscr{U}^*(y_n)(k)} R(k, p, y_n(\Lambda(k, p)))$, $n = 1, \ldots, T$, generate a sequence of actions analogous to the $\widehat{p}$ actions used in the constructions of the previous section. Without state variables, paths consisting of repeated play of $\widehat{p}$ followed by an optimal path for a "blocking player" were used to construct revision-proof strategies. With state variables the situation is similar except that the policy correspondences $\Psi(k) = \text{argmax}_{\mathscr{U}^*(y_n)(k)} R(k, p, y_n(\Lambda(k, p)))$ are used to construct play leading up to the blocking player. Let $w_{n+1}(k) = \max_{\Gamma(k)} Q(u(k,p), y_n(\Lambda(k,p)))$ and let $w_\infty(k) = \liminf_{n\to\infty} w_n(k)$. For an appropriate choice of n a payoff arbitrarily close to $w_\infty(k)$ can be attained by the first player in a sequence of the sort described above.

We impose the following analogue of Assumption 4. It implies that players along the continuation of a best path (for a blocking player) are strictly better off than if a path with payoff close to $w_\infty(k)$ is pursued.

Assumption 4′. *For each $k \in \mathscr{K}$, $Q(u(k, \overline{p}(k)), \overline{y}(k)) > w_\infty(k)$.*

We can now state an analogue of Proposition 5 for quasi-recursive games with state variables. The proof follows the same logic as the proof of Proposition 5 and is available upon request. It invokes Assumption 2′, Appendix E, a generalization of Assumption 2 for games with state variables, to ensure that the continuation payoff function Y is well defined.

Proposition 6. *Let Assumption 2′ in Appendix E and Assumption 4′ hold. Consider a state $k \in \mathscr{K}$. Let $\widetilde{P}$ be any path such that for all t, $\inf_t W(\widetilde{P}_t) > w_\infty(k)$. There is a revision-proof strategy σ that implements the path $\widetilde{P}$.*

The following example illustrates.

6.1 Example: Saving with quasi-hyperbolic preferences

A decision-maker has an initial endowment of capital k_1 and access to a linear technology for producing output from capital $q = Ak$, $A > 0$. In each period t she divides current output Ak_t into consumption p_t and future capital k_{t+1}. Let $\mathscr{K}, \mathscr{P} = \mathbb{R}_+$ denote sets of capitals and consumptions. Capital evolves according to a function $\Lambda : \mathbb{R} \times \mathscr{P} \to \mathbb{R}$, where:

$$\Lambda(k, p) = Ak - p.$$

For $k \in \mathscr{K}$, define $\Gamma(k) = \{p | p \in [0, Ak]\}$. Let $\Lambda^1 = \Lambda$ and, for $t \geq 1$, define $\Lambda^{t+1} : \mathscr{K} \times \mathscr{P}^{t+1} \to \mathbb{R}$ as in (9). $\Lambda^{t+1}(k, P^{t+1})$ gives capital for use in $t + 2$ given an initial capital k and prior consumption choices P^{t+1}. The set of feasible consumption paths given capital k, $\Pi(k)$, is then defined as in (10). Let $\beta \in (0, \min(1, A^{\gamma-1}))$ and $\gamma \in (0, 1)$. Each player's continuation payoff function is identified with:

$$Y(k, P) = \begin{cases} \sum_{t=1}^{\infty} \beta^{t-1} \frac{p_t^{1-\gamma}}{1-\gamma} & \text{if } P \in \Pi(k) \\ -\infty & \text{otherwise,} \end{cases}$$

where the restriction of Y to Graph Π satisfies $0 \leq Y(k, P) \leq \overline{b}k^{1-\gamma}$, $\overline{b} = \frac{A^{1-\gamma}}{1-\gamma}(1 - (\beta A^{1-\gamma})^{\frac{1}{\gamma}})^{-\gamma}$ and Y encodes infeasible plans outside of Graph Π with $-\infty$ values. Y satisfies the recursion:

$$Y(k, p, P') = R(k, p, Y(\Lambda(k, p), P')),$$

where $R(k, p, y) = u(k, p) + \beta y$ and $u(k, p) = \frac{p^{1-\gamma}}{1-\gamma}$ if $p \in \Gamma(k)$ and is $-\infty$ otherwise. The decision-maker's payoff in each period is:

$$W(k, p, P') = Q(u(k, p), Y(\Lambda(k, p), P')),$$

where: $Q(u(k, p), y) = u(k, p) + \beta\delta y$, $\delta \in (0, 1)$. The following proposition characterizes the set of sub-game perfect continuation payoffs in this case.

Proposition 7. *The sub-game perfect continuation payoff correspondence $\mathscr{Y}$ is given by:*

$$\mathscr{Y}(k) = \left[\underline{b}k^{1-\gamma}, \overline{b}k^{1-\gamma}\right].$$

The minimal sub-game perfect continuation payoff function $\underline{y}(k) = \underline{b}k^{1-\gamma}$ is the unique function *satisfying:*

$$\underline{b}k^{1-\gamma} = \frac{(Ak - k'(k))^{1-\gamma}}{1-\gamma} + \beta\underline{b}k'(k)^{1-\gamma},$$

where

$$k'(k) = argmax_{[0,Ak]} \frac{(Ak - k')^{1-\gamma}}{1-\gamma} + \beta\delta\underline{b}k'^{1-\gamma},$$

i.e. it is the unique Markov perfect value function.

The proof of the proposition is quite long and in the interests of space is omitted. It is available from the authors on request. In addition, defining $b_1 = \overline{b}$ and for each $n \in \mathcal{N}$,

$$b_{n+1}k^{1-\gamma} = \frac{(Ak - k_{n+1}(k))^{1-\gamma}}{1-\gamma} + \beta b_n k_{n+1}(k)^{1-\gamma},$$

where

$$k_{n+1}(k) = \text{argmax}_{[0,Ak]} \frac{(Ak - k')^{1-\gamma}}{1-\gamma} + \beta\delta b_n k'^{1-\gamma},$$

we have that $b_n \downarrow \underline{b}$. Thus, the value function obtained by having a terminal player obtain her optimal payoff and a sequence of preceding players choosing their best current payoff converges to the Markov (worst sub-game perfect) payoff function. Application of Proposition 6 implies that all paths with payoffs above the Markov (worst sub-game perfect) value function are revision-proof. In fact, in this case, a slight strengthening of the result establishes that all paths are revision-proof.

7 Revision-Proofness and the Literature

We relate revision-proofness to similar concepts in the literature.

7.1 Asheim's notion of revision-proofness

Asheim (1997) introduces an equilibrium concept he calls *revision-proofness* and we will call *A-revision-proofness*. We now show that despite apparent differences in their formulations, the concepts are in fact identical. We first introduce some preliminary concepts and definitions from Asheim. A strategy correspondence S associates each history P^{t-1} with a set of strategies, i.e. $S : \cup_{t\in\mathcal{N}} \mathscr{P}^{t-1} \to 2^{\mathscr{S}}$.

Definition 9. (Asheim)

1. A strategy correspondence is *internally stable* if for all $t \in \mathcal{N}$, $P^{t-1} \in \mathscr{P}^{t-1}$ and $\sigma \in S(P^{t-1})$, there is no $P^{t+r-1} = (P^{t-1}, P^r)$ and $\sigma' \in S(P^{t+r-1})$ such that $V_{t+r}(P^{t+r-1}, \sigma') > V_{t+r}(P^{t+r-1}, \sigma)$.

2. A strategy correspondence is *externally stable* if for all $t \in \mathcal{N}$, $P^{t-1} \in \mathscr{P}^{t-1}$, $\sigma' \in \mathscr{S} \backslash S(P^{t-1})$, there is some $P^{t+r-1} = (P^{t-1}, P^r)$ and $\sigma \in S(P^{t+r-1})$ such that $V_{t+r}(P^{t+r-1}, \sigma) > V_{t+r}(P^{t+r-1}, \sigma')$.

3. A strategy correspondence is stable if it is both internally and externally stable.

Definition 10. (Asheim) σ is *A-revision-proof* at P^{t-1} if it belongs to $S(P^{t-1})$, where S is stable strategy correspondence.

Lemma 2.

1) Let S be a strategy correspondence internally stable, then there is a function $\Psi : \cup_{t\in\mathcal{N}} \mathscr{P}^{t-1} \to \mathbb{R}$ such that for each P^{t-1} and all $\sigma \in S(P^{t-1})$, $V_t(P^{t-1}, \sigma) = \Psi(P^{t-1})$.

2) Let S be a stable strategy correspondence with payoff function Ψ. Let $\sigma \in S(P^{t-1})$, then for all $P^{t+r-1} = (P^{t-1}, P^r)$, $\sigma \in S(P^{t+r-1})$ and $V_{t+r}(P^{t+r-1}, \sigma) = \Psi(P^{t+r-1})$.

Proof. Part 1) If $\sigma, \sigma' \in S(P^{t-1})$, then immediately from the definition of internal stability $V_t(P^{t-1}, \sigma) \geq V_t(P^{t-1}, \sigma') \geq V_t(P^{t-1}, \sigma)$. Set $\Psi(P^{t-1})$ to this common value. Part 2) Let $P^{t+r-1} = (P^{t-1}, P^r)$. If $\sigma \in S(P^{t-1})$, then, either $\sigma \in S(P^{t+r-1})$, in which case the preceding step implies $V_{t+r}(P^{t+r-1}, \sigma) = \Psi(P^{t+r-1})$ or $\sigma \notin S(P^{t+r-1})$. In the latter case, by external stability of S, there is a successor history $P^{t+r+s-1} = (P^{t-1}, P^r, P^s) = (P^{t-1}, P^{r+s})$ such that $\Psi(P^{t+r+s-1}) > V_{t+r+s}(P^{t+r+s-1}, \sigma)$. But since $\sigma \in S(P^{t-1})$, internal stability implies that $V_{t+r+s}(P^{t+r+s-1}, \sigma) \geq \Psi(P^{t+r+s-1})$. We deduce that, in fact, if $\sigma \in S(P^{t-1})$, then $\sigma \in S(P^{t+r-1})$ and $V_{t+r}(P^{t+r-1}, \sigma) = \Psi(P^{t+r-1})$. □

The preceding lemma implies that if σ is A-revision-proof at P^{t-1}, then it is A-revision-proof at all successor histories P^{t+r-1}. We will say that σ is A-revision-proof if it is A-revision-proof at P^0 (and, hence, at all successor histories). If σ is A-revision proof, then all other strategies σ' must either deliver the same payoffs as σ after each history (and, hence, can be taken to belong to the same stable correspondence as σ) or they must deliver a strictly lower payoff than σ after some history. Consequently, we have the following result.

Proposition 8. *σ is revision-proof if and only if it is A-revision-proof.*

Proof. (If) Suppose that σ is A-revision-proof and that there is a history P^{t-1} and an alternative strategy σ' such that $\forall P^r$, $V_{t+r}(P^{t+r-1}, \sigma') \geq V_{t+r}(P^{t+r-1}, \sigma)$ with the inequality strict for at least one P^r. The A-revision-proofness of σ implies that there is a stable strategy correspondence S such that $\sigma \in S(P^0)$. Now, σ' cannot be in $S(P^{t-1})$ as well, since then S would fail to be internally stable. On the other hand if $\sigma' \notin S(P^{t-1})$, then S fails to be externally stable since for all P^{t+r-1}, $V_{t+r}(P^{t+r-1}, \sigma') \geq V_{t+r}(P^{t+r-1}, \sigma) = \Psi(P^{t+r-1})$, where Ψ is the payoff function associated with S and the equality stems from the previous lemma. Hence, if σ is A-revision-proof, it is revision-proof.

(Only If) Suppose that σ is revision-proof. For each P^{t-1}, define $S(P^{t-1})$ to be the set of strategies that give the same payoffs at P^{t-1} and all successor histories as σ. Each $S(P^{t-1})$ is non-empty since it contains σ. Since for all P^{t-1}, the members of $S(P^{t-1})$ give the same payoffs at P^{t-1} and all successor histories, S is internally stable. Revision-proofness and the definition of S imply that for all P^{t-1} any $\sigma' \notin S(P^{t-1})$ must give a strictly lower payoff after some successor history. Thus, S is externally stable. Since $\sigma \in S(P^0)$, σ is A-revision-proof. □

7.2 Kocherlakota's reconsideration-proofness

Kocherlakota (1996) specializes Farrell and Maskin (1989)'s definition of renegotiation-proofness to games played by dynasties. He considers only repeated games without history dependence. Player preferences are defined by a function $W : \mathscr{P}^\infty \to \mathbb{R}$ according to, for all t and P, $U(t, P) = W(P_t)$. The payoffs induced by a strategy σ are given by: $V_t(P^{t-1}, \sigma) = W(\Phi(\sigma | P^{t-1}))$.

Definition 11. (Kocherlakota) In a repeated game, a strategy σ is *reconsideration-proof* if (1) it is sub-game perfect, (2) it is payoff stationary, i.e. for all P^{t-1}, $V_t(P^{t-1}, \sigma) = \overline{V}$ for some number $\overline{V}$, (3) there is no other sub-game perfect equilibrium σ' such that for all P^{t-1}, $V_t(P^{t-1}, \sigma') = \overline{V}' > \overline{V}$.

Proposition 9. *If, in a repeated game, σ is revision-proof and payoff stationary, then it is reconsideration-proof.*

Proof. If σ is revision-proof, then it is also sub-game perfect and, hence, satisfies 1) of the reconsideration -proof definition. Payoff stationarity implies that σ satisfies 2) of the reconsideration-proof definition. Additionally, there is no alternative sub-game perfect equilibrium such that for all P^{t-1}, $V_t(P^{t-1}, \sigma') \geq V_t(P^{t-1}, \sigma) = V$ with strict inequality after some P^{t-1}. In particular 3) of the reconsideration-proof definition is satisfied. □

Revision-proofness ensures (1) of Definition 11 and if (2) holds, ensures (3) of this definition as well. However, revision-proofness, does not require (2). Reconsideration-proofness is, therefore, a refinement of revision-proofness in repeated games. Revision-proofness has the advantage that it can be applied to dynamic games, reconsideration-proofness cannot.

8 Conclusion

What types of outcomes can we expect to emerge when players with conflicting objectives sequentially take actions that affect their predecessors and, possibly, their successors? To address the question we consider strategies that are robust to joint deviations of successive players. We call these strategies revision-proof. Although this concept is not new, the literature contains very little prior characterization.

The main message of this paper is that if it is possible to support a path with off-equilibrium play that deters unilateral defection and blocks every infinite revision by eventually giving some player more than she can be obtain from joining the revision, then the path is revision-proof. In games where there is some consensus about how to play or not play, this may only be possible for a small set of paths and the concept may be quite selective. In quasi-recursive games where optimal actions for successors are quite costly for predecessors, but optimal actions for predecessors are not too costly for successors, the concept is quite permissive with respect to paths.

We have given applications of our results to games played by overlapping generations of players, to savings games played by an agent with quasi-hyperbolic discounting preferences and to a simple macroeconomic policy game. In the latter case, the macroeconomic policy game mapped directly into a reduced form quasi-recursive game. For macroeconomic policy games in which private households take actions that influence the physical state of a later policymaker, a mapping into a quasi-recursive game exists, but it is more complicated. In these settings, the associated reduced form quasi-recursive game explicitly incorporates private households' beliefs about future policy actions, as well as the actions themselves. The revision-proofness concept can be extended to these games and we conjecture that analogues of our results are available for them. We also conjecture that our procedure for designing revision-proof strategies, with off-equilibrium payoff build ups culminating in blocking players, has application to a broader range of quasi-recursive and non-quasi recursive games. We leave these extensions to future work.

Appendix

A Proofs for Sections 3 and 4

Proof of Lemma 1. (Necessity) Suppose that σ is revision-proof for $\mathscr{G}(U)$ and that at some P^{t-1} there is a successful revision path $P' = \{p'_r\}_{r=1}^{\infty}$. Modify σ to obtain a new strategy σ' such that $\Phi(\sigma'|P^{t-1}) = P'$, but otherwise σ' equals σ. Then for all r, $V_{t+r}(P^{t-1}, P'^r, \sigma') \geq V_{t+r}(P^{t-1}, P'^r, \sigma)$ with strict inequality for at least one r. Also, for all dates r and histories $P^r \neq P'^r$, $\Phi(\sigma'|P^{t-1}, P^r) = \Phi(\sigma|P^{t-1}, P^r)$, so that $V_{t+r}(P^{t-1}, P^r, \sigma')$ $= V_{t+r}(P^{t-1}, P^r, \sigma)$. This contradicts the assumed revision-proofness of σ. Hence, at no P^{t-1} is there a successful revision path P'.

(Sufficiency) Let σ be such that there are no successful revision paths at any history P^{t-1}. Then, by Assumption 1, $\sigma \in \mathscr{S}(U)$. Let P^{t-1} and σ' be arbitrary. Let P^s be a history such that:

$$V_{t+s}(P^{t-1}, P^s, \sigma') > V_{t+s}(P^{t-1}, P^s, \sigma).$$

Thus, $U(t+s, P^{t-1}, P^s, \Phi(\sigma'|P^{t-1}, P^s)) > V_{t+s}(P^{t-1}, P^s, \sigma)$. Since there are no successful revision paths at any history, it follows that there is some r and sub-history P^r of $(P^s, \Phi(\sigma'|P^{t-1}, P^s))$ such that:

$$U(t+r, P^{t-1}, P^s, \Phi(\sigma'|P^{t-1}, P^s)) < V_{t+r}(P^{t-1}, P^r, \sigma).$$

Hence, the adoption of σ' at P^{t-1} either leaves all payoffs unaltered or strictly lowers some. Since σ' and P^{t-1} were arbitrary, it follows that σ is revision-proof for $\mathscr{G}(U)$. □

Proof of Proposition 4. Sufficiency is obvious. We remark that by Assumption 1, each $\overline{U}_t(P^{t-1})$ is greater than $-\infty$ and so a strategy σ attaining the payoffs $\overline{U}_t(P^{t-1})$ is in $\mathscr{S}(U)$. We turn to necessity. Suppose that σ is sub-game perfect (and so in $\mathscr{S}(U)$) and that there is some t and P^{t-1} such that $\overline{U}(t, P^{t-1}) > V_t(P^{t-1}, \sigma)$. Then there is a path P' such that $(P^{t-1}, P') \in \text{Dom}\, \Pi_t(P^{t-1})$ and $U(t, P^{t-1}, P') > V_t(P^{t-1}, \sigma) + \varepsilon$ for some small $\varepsilon > 0$. Continuity in the product topology implies a T such that $U(t, P^{t-1}, P'^T, \Phi(\sigma|P^{t-1}, P'^T)) + \varepsilon > U(t, P^{t-1}, P')$. So, $U(t, P^{t-1}, P'^T, \Phi(\sigma|P^{t-1}, P'^T)) > V_t(P^{t-1}, \sigma)$. Consider the player at $t+T-1$. By sub-game perfection, $U(t+T-1, P^{t-1}, P'^{T-1}, \Phi(\sigma|P^{t-1}, P'^{T-1})) \geq U(t+T-1, P^{t-1}, P'^T, \Phi(\sigma|P^{t-1}, P'^T))$. Time consistency then implies:

$$U(t, P^{t-1}, P'^{T-1}, \Phi(\sigma|P^{t-1}, P'^{T-1})) \geq U(t, P^{t-1}, P'^T, \Phi(\sigma|P^{t-1}, P'^T)) > V(t, P^{t-1}, \sigma).$$

Iterating back from T in this way, we obtain: $V_t(P^{t-1}, \sigma) = U(t, P^{t-1}, \Phi(\sigma|P^{t-1})) > V_t(P^{t-1}, \sigma)$. This is a contradiction. We conclude that $\overline{U}(t, P^{t-1}) = V_t(P^{t-1}, \sigma)$. □

B Sub-game perfection with quasi-recursive payoffs

Let $\mathscr{C}$ denote the set of bounded, continuous functions $f : \mathscr{P}^\infty \to \mathbb{R}$ equipped with the sup-norm $\|\cdot\| : \mathscr{C} \to \mathbb{R}_+$, with for $f \in \mathscr{C}$, $\|f\| = \sup_{\mathscr{P}^\infty} |f(P)|$, and the partial order

$\geq$, with for $f, f' \in \mathscr{C}$, $f \geq f'$ if for all $P \in \mathscr{P}^\infty$, $f(P) \geq f'(P)$. Let $\mathbf{0}$ denote the zero function: $\mathbf{0} : \mathscr{P}^\infty \to 0$.

Proposition 10. *Suppose Assumption 2 holds. Then there is a unique function $Y \in \mathscr{C}$ such that $Y(p, P') = R(p, Y(P'))$. Also,* $\lim \|T^n(\mathbf{0}) - Y\| \to 0$.

Proof. Define the operator T on the domain $\mathscr{C}$ by $T(f)(p, P) = R(p, f(P))$. Since R is increasing in its second argument, T is increasing. Also, $|T(\mathbf{0})| = |R(p, 0)| < \infty$. Finally, by the discounting property of R in Assumption 2,

$$T(f + b)(p, P) = R(p, f(P) + b) \leq R(p, f(P)) + \beta b = T(f)(p, P) + \beta b$$

and so T satisfies a discounting property as well. This verifies the assumptions of Becker and Boyd (1997)'s generalization of Blackwell's theorem. It follows that $T : \mathscr{C} \to \mathscr{C}$ and has a unique fixed point in $\mathscr{C}$. Also, for any P, $|Y(P) - T^n(\mathbf{0})(P)| \leq \beta^n |Y(P_n)| \leq \beta^n \|Y\|$. Since the last term converges to 0 as n converges to ∞, we have $\lim \|T^n(\mathbf{0}) - Y\| \to 0$. $\square$

Let $\mathscr{Y}_0 := [-\|Y\|, \|Y\|]$. Evidently, $Y(\mathscr{P}^\infty) \subset \mathscr{Y}_0 := [-\|Y\|, \|Y\|]$. Given Assumption 2, $W(p, P) = Q(u(p), Y(P))$ is also continuous with respect to the relative product topology on $\mathscr{P}^\infty$. Define the $\mathscr{B}$-operator as follows.

Definition 12. Let $\mathscr{B} : 2^{\mathbb{R}} \to 2^{\mathbb{R}}$ be given by for all $\mathscr{Y}' \subset \mathbb{R}$,

$$\mathscr{B}(\mathscr{Y}') = \left\{ y \,\middle|\, \exists (p, y') \in \mathscr{P} \times \mathscr{Y}', \text{ with } y = R(p, y'),\ Q(u(p), y') \geq \sup_{\mathscr{P}} \inf_{\mathscr{Y}'} Q(u(p''), y'') \right\}.$$

Let $\mathscr{Y} = \{y | y = Y(\Phi(\sigma)) \text{ with } \sigma \text{ sub-game perfect}\}$.

Proposition 11. *(i) $\mathscr{Y} = \mathscr{B}(\mathscr{Y})$. (ii) If $\mathscr{W} \subset \mathbb{R}$ satisfies $\mathscr{W} \subseteq \mathscr{B}(\mathscr{W})$, then $\mathscr{W} \subseteq \mathscr{Y}$. (iii) $\mathscr{Y}$ is a compact interval, $[\underline{y}, \overline{y}]$.*

Proof. (i) and (ii) are immediate applications of arguments in Abreu, Pearce, and Stacchetti (1990) and are omitted. Endow $2^{\mathbb{R}}$ with the set inclusion ordering and define set convergence in the Kuratowski sense. For monotone, decreasing sequences of closed sets $\{C_n\}$, $C_n \subset \mathbb{R}$, the (Kuratowski) limit exists and is given by $\lim_{n\to\infty} C_n := \cap_{n=1}^\infty C_n$. $\mathscr{B}$ is monotone in the set inclusion ordering. In addition, $\mathscr{B}(\mathscr{Y}_0) \subset \mathscr{Y}_0$. Thus, the sequence $\{\mathscr{B}^n(\mathscr{Y}_0)\}$ is monotone decreasing.

Let $\mathscr{W}$ be a compact interval. Let y_1, y_2 belong to $\mathscr{B}(\mathscr{W})$. Then there exists (p_1, y_1') and (p_2, y_2') each in $\mathscr{P} \times \mathscr{W}$ such that for $k = 1, 2$, $y_k = R(p_k, y_k')$ and $Q(u(p_k), y_k') \geq \sup_{\mathscr{P}} \inf_{\mathscr{W}} Q(u(p), y)$. Let $H : [0, 1] \to \mathbb{R}$ be defined as for all $\psi \in [0, 1]$, $H(\psi) = R(\psi p_1 + (1 - \psi) p_2, \psi y_1' + (1 - \psi) y_2')$. H is well defined by the convexity of $\mathscr{P}$ and $\mathscr{W}$. Let $\lambda \in [0, 1]$, then by continuity of R and, hence, H, there is a ψ such that $(1 - \lambda) y_1 + \lambda y_2 = R(\psi p_1 + (1 - \psi) p_2, \psi y_1' + (1 - \psi) y_2')$. By concavity of u and Q, $Q(u(\psi p_1 + (1 - \psi) p_2), \psi y_1' + (1 - \psi) y_2') \geq \psi Q(u(p_1), y_1') + (1 - \psi) Q(u(p_2), y_2') \geq \sup_{\mathscr{P}} \inf_{\mathscr{W}} Q(u(p), y)$. Hence, $\lambda y_1 + (1 - \lambda) y_2 \in B(\mathscr{W})$ and $B(\mathscr{W})$ is an interval. Let $y_\infty \in \text{cl}\,(B(\mathscr{W}))$. Then there is a sequence $\{y_k\}$ with each $y_k \in B(\mathscr{W})$ and $\lim_{k\to\infty} y_k = y_\infty$. For each y_k there is a

$(p_k, y'_k) \in \mathscr{P} \times \mathscr{W}$ such that $y_k = R(p_k, y'_k)$ and $Q(u(p_k), y'_k) \geq \sup_{\mathscr{P}} \inf_{\mathscr{W}} Q(u(p), y)$. Since $\mathscr{P}$ and $\mathscr{W}$ are compact, the sequence $\{p_k, y'_k\}$ admits a convergent subsequence (p_{k_n}, y'_{k_n}) with limit (p_∞, y'_∞). By the continuity of R, u and Q, $y_\infty = \lim R(p_{k_n}, y'_{k_n}) = R(p_\infty, y'_\infty)$ and each $Q(u(p_{k_n}), y'_{k_n}) \geq \sup_{\mathscr{P}} \inf_{\mathscr{W}} Q(u(p), y')$ so $Q(u(p_\infty), y'_\infty) \geq \sup_{\mathscr{P}} \inf_{\mathscr{W}} Q(u(p), y')$. Hence, $y_\infty \in B(\mathscr{W})$ and so $B(\mathscr{W})$ is closed. Since $\mathscr{P}$ and $\mathscr{W}$ are compact and R is continuous, $B(\mathscr{W})$ is bounded. Combining results, $\mathscr{B}(\mathscr{W})$ is a compact interval.

Hence, since $\mathscr{Y}_0$ is a compact interval, the $\{\mathscr{B}^n(\mathscr{Y}_0)\}$ is a decreasing sequence of compact intervals. Thus, $\mathscr{Y}_\infty := \lim_{n\to\infty} \mathscr{B}^n(\mathscr{Y}_0) = \cap_{n=0}^{\infty}[\underline{y}_n, \overline{y}_n] = [\underline{y}, \overline{y}]$, where $[\underline{y}_n, \overline{y}_n] = \mathscr{B}^n(\mathscr{Y}_0)$, $\underline{y} = \sup_n \underline{y}_n$ and $\overline{y} = \inf_n \overline{y}_n$. Since $\mathscr{Y} \subseteq \mathscr{Y}_0$, $\mathscr{B}(\mathscr{Y}_0) \subseteq \mathscr{Y}_0$, $\mathscr{Y} = \mathscr{B}(\mathscr{Y})$ and $\mathscr{B}$ is monotone, we have for all n, $\mathscr{Y} \subseteq \mathscr{B}^n(\mathscr{Y}_0)$. Thus, $\mathscr{Y} \subseteq \mathscr{Y}_\infty$. For the reverse inclusion see Abreu, Pearce, and Stacchetti (1990). This proves (iii). □

By Proposition 11 (iii), it is sufficient to characterize the endpoints $\underline{y}$ and $\overline{y}$ of $\mathscr{Y}$. By Proposition 11 (i), these endpoint payoffs satisfy:

$$\underline{y} = \min \left\{ R(p, y') \;\middle|\; \text{s.t. } (p, y') \in \mathscr{P} \times [\underline{y}, \overline{y}] \text{ and } Q(u(p), y') \geq Q(u^*, \underline{y}) \right\} \tag{MIN}$$

and

$$\overline{y} = \max \left\{ R(p, y') \;\middle|\; \text{s.t.} (p, y') \in \mathscr{P} \times [\underline{y}, \overline{y}] \text{ and } Q(u(p), y') \geq Q(u^*, \underline{y}) \right\}, \tag{MAX}$$

where $u^* = \max_{\mathscr{P}} u(p)$. In addition, by Proposition 11 (ii) if $\underline{y}'$ and $\overline{y}'$ satisfy the above conditions, then $\underline{y} \leq \underline{y}' \leq \overline{y}' \leq \overline{y}$. Let: $\mathscr{X} = \{p \in \mathscr{P} | Q(u(p), \overline{y}) \geq Q(u^*, \underline{y})\}$.

Lemma 3.

1. *The solution to* (MIN) *is attained by a pair* $(\underline{p}, y'(\underline{p}))$*, where:*

$$\underline{p} \in \underset{\mathscr{X}}{\operatorname{argmin}}\, R(p, y'(p))$$

and for each $p \in \mathscr{X}$*,* $y'(p)$ *is the unique element of* $\mathscr{Y}$ *satisfying* $Q(u(p), y'(p)) = Q(u^*, \underline{y})$*. Either* $y'(\underline{p}) = \underline{y}$ *in which case* $\underline{y} = R(\underline{p}, \underline{y})$ *and* $u(\underline{p}) = u^*$ *or* $y'(\underline{p}) > \underline{y}$ *in which case* $\underline{y} > Y(\underline{p}, \underline{p}, \ldots)$*.*

2. *There is a* $\overline{p} \in \mathscr{X}$ *such that* $\overline{y} = R(\overline{p}, \overline{y})$*. If* $Q(u(\overline{p}), \overline{y}) > Q(u^*, \underline{y})$*, then* $\overline{y} = Y(\overline{p}, \overline{p}, \ldots) = \max_{\mathscr{P}^\infty} Y(P)$*.*

Proof. As previously noted, (MIN) has a solution. Let $(p, y') \in \mathscr{P} \times \mathscr{Y}$ be such a solution. Feasibility of (p, y') implies $Q(u(p), y') \geq \sup_{\mathscr{P}} Q(u(p'), \underline{y})$. Suppose that $y' > \underline{y}$ and $Q(u(p), y') > Q(u^*, \underline{y})$. Then, by the monotonicity of $R(p, \cdot)$ and the continuity of $Q(u(p), \cdot)$, the payoff $R(p, y')$ can be weakly reduced by lowering y'. Thus, there must be a solution to (MIN) with either $Q(u(p), y') = Q(u^*, \underline{y})$ or $y' = \underline{y}$. However, if $y' = \underline{y}$, then $Q(u^*, \underline{y}) \geq Q(u(p), y')$ and feasibility of (p, y') implies that $Q(u(p), y') = Q(u^*, \underline{y})$

for this case as well. For each p in $\mathscr{X}$, by the continuity and monotonicity of $Q(u(p),\cdot)$, there is a unique $y'(p) \in \mathscr{Y}$ such that: $Q(u(p),y'(p)) = Q(u^*,\underline{y})$. Thus, wlg (MIN) may be replaced with the restricted problem: $\min_{\mathscr{X}} R(p,y'(p))$, which has a solution since (MIN) does. Let $\underline{p}$ denote such a solution. If $y'(\underline{p}) = \underline{y}$, then we have immediately that: $\underline{y} = R(\underline{p},\underline{y})$ and $Q(u(\underline{p}),\underline{y}) = Q(u^*,\underline{y})$. Hence, $u(\underline{p}) = u^*$ and the minimal payoff and minimal continuation payoff are attained by repetition of $\underline{p}$. On the other hand, if $y'(\underline{p}) > \underline{y}$, then

$$\underline{y} = R(\underline{p},y'(p)) > R(\underline{p},\underline{y}) = R(\underline{p},R(\underline{p},y'(\underline{p}))) > R(\underline{p},R(\underline{p},\underline{y})) > \ldots > Y(\underline{p},\underline{p},\ldots)$$

and the minimal continuation payoff exceeds that from repetition of $\underline{p}$.

For part 2 of the Lemma, let $(\overline{p},y)$ be a solution to (MAX). Since R is non-decreasing in its second argument, $(\overline{p},\overline{y})$ is a solution as well and $\overline{y} = R(\overline{p},\overline{y})$. Since $\mathscr{P}$ is compact, $\mathscr{P}^\infty$ is compact in the product topology. The Lipschitz condition on R ensures that Y is continuous in this topology and so, by Weierstrass's theorem, there is a $P' = (p',P'')$ that maximizes Y on $\mathscr{P}^\infty$ and attains a maximal payoff $\overline{\mu}$. Since R non-decreasing in its second argument it follows that P'' attains $\overline{\mu}$ as well, i.e. $\overline{\mu} = R(p',\overline{\mu})$. Suppose that $Q(u(\overline{p}),\overline{y}) > Q(u^*,\underline{y})$ and $\overline{\mu} = R(p',\overline{\mu}) > \overline{y} = R(\overline{p},\overline{y})$. Let $\lambda \in [0,1]$, $p^\lambda = \lambda\overline{p} + (1-\lambda)p'$ and $y^\lambda = \lambda\overline{y} + (1-\lambda)\overline{\mu}$. Then, by the concavity of R, for all $\lambda \in (0,1]$, $R(p^\lambda,y^\lambda) \geq \lambda R(\overline{p},\overline{y}) + (1-\lambda)R(p',\overline{\mu}) > R(\overline{p},\overline{y})$. By the continuity of u and Q, there is a $\lambda \in (0,1]$ small enough that, $Q(u(p^\lambda),y^\lambda) \geq Q(u^*,\underline{y})$. Now, $R(p^\lambda,y^\lambda) \geq \lambda R(\overline{p},\overline{y}) + (1-\lambda)R(p',\overline{\mu}) = y^\lambda$. It follows using the monotonicity property of R and definition of Y that $y^\psi := Y(p^\lambda,p^\lambda,\ldots) \geq y^\lambda > \overline{y}$. But then, since Q is increasing in its second argument, $Q(u(p^\lambda),y^\psi) > Q(u(p^\lambda),y^\lambda) \geq Q(u^*,\underline{y})$. But this contradicts Proposition 11 (ii). □

C Proof of Proposition 5

Proof. Throughout this proof, we repeatedly use the continuity, monotonicity and discounting properties of u, Q and R as described in Assumption 2 and, in particular, their implication that W is continuous. We also use Assumption 3, $\mathscr{R}^*$ does not depend on y. Let $\widehat{p} \in \mathscr{R}^* \subset \mathscr{U}^*$ and $\overline{p}$ be as in the preceding text. Let $\overline{\mathscr{N}} = \{0,\infty\} \cup \mathscr{N}$. Define $P : \overline{\mathscr{N}} \to \mathscr{P}^\infty$ as follows. For each $T \in \{0\} \cup \mathscr{N}$, let $P(T) = (\widehat{p},\ldots,\widehat{p},\overline{p},\overline{p},\ldots)$, with T the number of periods until the sequence $(\overline{p},\overline{p},\ldots)$ begins. Let $P(\infty) = \widehat{P} = (\widehat{p},\widehat{p},\ldots)$. Choose $\tau : \mathscr{P} \to \mathscr{N}$ to satisfy for each $p \in \mathscr{P}\backslash\mathscr{U}^*$, $\widehat{w} > W(p,P(\tau(p)))$, which is possible given the definition of $\widehat{w}$ and $\mathscr{U}^*$ and the continuity of W.

Set σ so that $\Phi(\sigma) = \widetilde{P}$. For each $t \in \mathscr{N}$, if $\Phi(\sigma|\widetilde{P}^{t-1}) \neq P(T)$, $T \in \overline{\mathscr{N}}$ and $p \neq \widetilde{p}_t$, then set $\Phi(\sigma|\widetilde{P}^{t-1},p) = P(\infty)$. For all $t \in \mathscr{N}$ and P^{t-1}, (i) if $\Phi(\sigma|P^{t-1}) = P(\infty)$, then for $p \in \mathscr{P}\backslash\mathscr{U}^*$, set $\Phi(\sigma|P^{t-1},p) = P(\tau(p))$ and for $p \in \mathscr{U}^*$, set $\Phi(\sigma|P^{t-1},p) = P(\infty)$, (ii) if $\Phi(\sigma|P^{t-1}) = P(T)$, $T \in \mathscr{N}$, then for each $p \in \mathscr{P}$, set $\Phi(\sigma|P^{t-1},p) = P(T-1)$, (iii) if $\Phi(\sigma|P^{t-1}) = P(0)$, and $p \neq \overline{p}$, then set $\Phi(\sigma|P^{t-1},p) = P(\infty)$.

It is easy to see that σ is sub-game perfect. We now show that it is revision-proof. By Lemma 1, it is sufficient to show that there is no history P^t, $t \in \{0\} \cup \mathscr{N}$, and revision path $P'' = \{p''_r\}_{r=1}^\infty$ with histories P''^r and continuations P''_r such that for $r \in \mathscr{N}$,

$W(P''_r) \geq W(\Phi(\sigma|P^t, P''^{r-1}))$ with at least one of these inequalities strict. Suppose that such a history P^t and revision path P'' exists, we seek a contradiction. For each $r \in \mathcal{N}$, let $W_r := W(P''_r)$ and $Y_r := Y(P''_r)$. It is immediate that $P'' \neq \Phi(\sigma|P^t)$. We distinguish three further cases. Case 1: $\Phi(\sigma|P^t) = P(\infty)$ and for all $r \in \mathcal{N}$, $p''_r \in \mathscr{U}^*$. Case 2: either $\Phi(\sigma|P^t) = \widetilde{P}_{t+1} \neq P(T)$, $T \in \overline{\mathcal{N}}$, or $\Phi(\sigma|P^t) = P(0)$, there is a first date $r_0 \in \mathcal{N}$ at which $p''_{r_0} \neq \sigma_{t+r_0}(P^t, P''^{r_0-1})$ and $\Phi(\sigma|P^t, P''^{r_0}) = P(\infty)$ and from $r_0 + 1$ onwards each $p''_r \in \mathscr{U}^*$. Case 3: There is a first date r_0 at which $\Phi(\sigma|P^t, P''^{r_0-1}) = P(T)$, $T \in \mathcal{N}$. It is easily checked that all P^t and P'' with $P'' \neq \Phi(\sigma|P^t)$ fall into one of these cases.

Suppose that P'' belongs to Case 1. Since in this case every p''_r belongs to $\mathscr{U}^*$, the definition of σ implies that for all $r \in \mathcal{N}$, $\Phi(\sigma|P^t, P''^{r-1}) = P(\infty)$ and $W(\Phi(\sigma|P^t, P''^{r-1})) = \widehat{w}$. On the other hand since some p''_r may belong to $\mathscr{U}^* \backslash \mathscr{R}^*$ and, therefore, may not be maximal for R on $\mathscr{U}^*$, we have that $W_r = Q(u(p''_r), Y(P''_{r+1})) \leq Q(u(p''_r), \widehat{y}) = \widehat{w}$. Hence, there is no player for whom $W_r > W(\Phi(\sigma|P^t, P''^{r-1})$. But this contradicts the assumed property of P'' and so, in fact, P'' cannot belong to Case 1.

Next suppose that P'' belongs to Case 2. Similar to Case 1, for $r > r_0$, $W_r \leq \widehat{w} = W(\Phi(\sigma|P^t, P''^{r-1}))$. By (i) in the proposition and Assumption 4, $\widehat{w} \leq W(\Phi(\sigma|P^t, P''^{r_0-1}))$. Also, since p''_{r_0} need not be in $\mathscr{U}^*$ and some later p''_r values need not be in $\mathscr{R}^*$, $W_{r_0} = Q(u(p''_{r_0}), Y(P''_{r_0+1})) \leq Q(u(p''_{r_0}), \widehat{y}) \leq \widehat{w}$. In addition, by (ii) in the proposition and Assumption 4, either the player at $t + r_0$ is strictly worse off adhering to P'', $W_{r_0} < W(\Phi(\sigma|P^t, P''^{r_0-1}))$, or she is no worse off, $W_{r_0} = \widehat{w} = W(\Phi(\sigma|P^t, P''^{r_0-1}))$, and, if $t + r_0 > 1$, the player at $t + r_0 - 1$ receives a weakly lower continuation payoff from P'', $Y_{r_0} \leq Y(\Phi(\sigma|P^t, P''^{r_0-1}))$. Since P'' coincides with $\Phi(\sigma|P^t)$ between 1 and $r_0 - 1$, it follows that all players between to t and $t + r_0 - 1$ are weakly worse off adhering to P''. But then there is no player for whom $W_r > W(\Phi(\sigma|P^t, P''^{r-1})$ and so, in fact, P'' cannot belong to Case 2.

Finally, suppose that P'' belongs to Case 3. Then there is a first date $t + r_0 + 1$ at which $\Phi(\sigma|P^t, P''^{r_0}) = P(T)$ for some $T \in \mathcal{N}$. By Lemma 4 below, for $r = 1, \ldots, T$, $W(P''_{r_0+r}) = W(\Phi(\sigma|P^t, P''^{r_0+r-1}))$ and, if $t + r_0 > 0$, $Y_{r_0+1} \leq Y(\Phi(\sigma|P^t, P''^{r_0}))$. There are two sub-cases. In sub-case 3A, $r_0 > 0$, $\Phi(\sigma|P^t, P''^{r_0-1}) = P(\infty)$, $p''_{r_0} \notin \mathscr{P} \backslash \mathscr{U}^*$ and $\Phi(\sigma|P^t, P''^{r_0})$ $= P(\tau(p''_{r_0}))$. In this case, using the definition of $\tau(p''_{r_0})$, Lemma 4 and the monotonicity of Q, the player at $t + r_0$ is made strictly worse off since: $W(\Phi(\sigma|P^t, P''^{r_0-1})) = \widehat{w} > W(p''_{r_0}, P(\tau(p''_{r_0}))) = Q(u(p''_{r_0}), Y(P(\tau(p''_{r_0}))) \geq Q(u(p''_{r_0}), Y_{r_0+1})$. But then P'' cannot belong to sub-case 3A. In sub-case 3B, $r_0 = 1$ and $\Phi(\sigma|P^t) = P(T)$. By Lemma 4, for $r = 1, \ldots, T$, $W(P''_r) = W(\Phi(\sigma|P^t, P''^{r-1}))$. If $P''_{T+1} = P(0) = (\overline{p}, \overline{p}, \ldots)$, i.e. there are no further defections from the strategy, then for all $r \in \mathcal{N}$, $W(P''_r) = W(\Phi(\sigma|P^t, P''^{r-1}))$ and there is no player for whom $W(P''_r) > W(\Phi(\sigma|P^t, P''^{r-1}))$. If there is a first $r_1 > T$ such that $p''_{r_1} \neq \overline{p}$ and for all $r > r_1$, $p''_r \in \mathscr{U}^*$, then following our analysis of Case 2, for all r, $W(P''_r) \leq W(\Phi(\sigma|P^t, P''^{r-1}))$ and again there is no player for whom $W(P''_r) > W(\Phi(\sigma|P^t, P''^{r-1}))$. Finally, if there is a first $r_1 > T$ such that $p''_{r_1} \neq \overline{p}$ and a first $r_2 > r_1$ such that $p''_r \in \mathscr{P} \backslash \mathscr{U}^*$, then following our analysis of sub-case 3A some player is made worse off. Thus, either no player is made strictly better off or some player is made strictly worse off and P'' cannot belong to sub-case 3B. This exhausts all possible cases and so there is no history P^t and revision path P'' such that latter weakly raises the payoff of all players relative to reversion to the strategy and strictly raises the payoff of some. It follows from Lemma 1 that the strategy is revision-proof. □

Lemma 4. *Let σ, $P : \{0, \infty\} \cup \mathscr{N} \to \mathscr{P}^\infty$ and P'' be as in the previous proof. For any $t \in \{0\} \cup \mathscr{N}$ and P^t, if $\Phi(\sigma|P^t) = P(T)$, $T \in \mathscr{N}$, then for $r = 1, \ldots, T$, $W(P''_r) = W(\Phi(\sigma|P^t, P''^{r-1}))$ and $Y(P'') \leq Y(\Phi(\sigma|P^t))$.*

Proof. As in the proof of Proposition 5, for all $r \in \mathscr{N}$, let $W_r = W(P''_r)$ and $Y_r = Y(P''_r)$. Consider play at $t + T$. Given that $\Phi(\sigma|P^t) = P(T)$, the construction of σ implies: $\Phi(\sigma|P^t, P''^{T-1}) = P(1)$. Thus, σ implements the best sub-game perfect path, $P(1) = (\widehat{p}, \overline{p}, \overline{p}, \ldots)$, for the $t + T$-th player. Since P'' at least matches the payoffs received under the strategy, it follows that:

$$W_T \geq W(P(1)) = Q(u(\widehat{p}), Y(P(0))) = Q(u^*, \overline{y}).$$

We claim that $W_T = W(P(1))$ and $Y_T \leq Y(P(1))$. Suppose the claim is false, then either A) $W_T > W(P(1))$ or B) $W_T = W(P(1))$ and $Y_T > Y(P(1))$. We show that in either of these cases, $Y_{T+1} > \overline{y}$ and P''_{T+1} is a sub-game perfect path. This contradiction establishes the claim.

Suppose A) $W_T > W(P(1))$, then $Q(u(p''_T), Y_{T+1}) = W_T > W(P(1)) = Q(u^*, \overline{y})$ and, since $u^* \geq u(p''_T)$ and Q is monotone, $Y_{T+1} > \overline{y}$. Suppose B) $W_T = W(P(1))$ and $Y_T > Y(P(1))$, then $R(p''_T, Y_{T+1}) = Y_T > Y(P(1)) = R(\widehat{p}, \overline{y})$ and since R is weakly monotone in its second argument, either $Y_{T+1} > \overline{y}$ or $p''_T \neq \widehat{p}$. If $Y_{T+1} \leq \overline{y}$ and $p''_T \neq \widehat{p}$, then, by definition of $\widehat{p}$, $p''_T \notin \mathscr{U}^*$ and $u^* = u(\widehat{p}) > u(p''_T)$. However, $W_T = Q(u(p''_T), Y_{T+1}) = Q(u^*, \overline{y})$ and so either $p''_T \in \mathscr{U}^*$ or $Y_{T+1} > \overline{y}$ must occur. Hence, in fact, $Y_{T+1} > \overline{y}$.

Suppose that for some $k \in \mathscr{N}$, $Y_{T+k} > \overline{y}$. Since:

$$\overline{y} = \max_{(p,y')} \{R(p, y'), y' \in [\underline{y}, \overline{y}], Q(u(p), y') \geq Q(u^*, \underline{y})\} \tag{14}$$

and $R(p''_{T+k}, Y_{T+k+1}) = Y_{T+k} > \overline{y}$, (p''_{T+k+1}, Y_{T+k+1}) is not feasible for (14). Either 1) $Q(u(p''_{T+k}), Y_{T+k+1}) < Q(u^*, \underline{y})$, 2) $Y_{T+k+1} < \underline{y}$ or 3) $Y_{T+k+1} > \overline{y}$. If 1) holds then the player at $t + T + k$ is strictly worse off under P'' than under the continuation of σ (which is sub-game perfect and, thus, gives a payoff weakly above $Q(u^*, \underline{y})$ after all histories). But this contradicts the assumed property of P''. So 1) cannot hold. If 2), but not 1) holds, then $R(p''_{T+k}, \underline{y}) \geq R(p''_{T+k}, Y_{T+k+1}) = Y_{T+k} > \overline{y}$. But since $(p''_{T+k}, \underline{y})$ is clearly feasible for (14), this contradicts the optimality of $\overline{y}$ in (14). Hence, 3) holds and $Y_{T+k+1} > \overline{y}$. By induction for all $k \in \mathscr{N}$, $Y_{T+k} = R(p''_{T+k}, Y_{T+k+1}) > \overline{y}$ with each $Q(u(p''_{T+k}), Y_{T+k+1}) \geq Q(u^*, \underline{y})$. But then, in fact, P''_{T+1} is a sub-game perfect path with continuation payoff $Y_{T+1} > \overline{y}$. This is a contradiction. We conclude that $W_T = W(P(1)) = W(\Phi(\sigma|P''^{T-1}))$ and $Y_T \leq Y(P(1))$.

At $t + T - 1$ after (P^t, P''^{T-2}), σ implements the path $P(2) = (\widehat{p}, \widehat{p}, \overline{p}, \overline{p}, \ldots)$. The $t + T - 1$-th player receives $Q(u^*, Y(P(1)))$ and since $u(p''_{T-1}) \leq u^*$ and $Y_T \leq Y(P(1))$, we have: $W_{T-1} = Q(u(p''_{T-1}), Y_T) \leq Q(u^*, Y(P(1))) = W(\Phi(\sigma|P''^{T-2}))$. However, by the assumed property of P'', $W_{T-1} \geq W(\Phi(\sigma|P''^{T-2}))$. So, $W_{T-1} = W(\Phi(\sigma|P''^{T-2}))$, $u(p''_{T-1}) = u(\widehat{p}) = u^*$ and $Y_T = Y(P(1))$. Moreover, $Y_{T-1} = R(p''_{T-1}, Y_T) \leq R(\widehat{p}, Y(P(1)))$

$= Y(P(2))$ since $Y_T = Y(P(1))$, $p''_{T-1} \in \mathscr{U}^*$ and $\widehat{p}$ is maximal for $R(p, Y(P(1)))$ on $\mathscr{U}^*$. Continuing with this logic back to $t+1$, we find that P'' must give each player between $t+1$ and $t+T$ current and continuation payoffs equal to those obtained under σ and the player at t a continuation payoff less than or equal to that obtained under σ, i.e. less than or equal to $Y(P(T))$. This completes the proof. $\square$

Corollary 2. *Let Assumptions 2 to 4 hold. Let $\widetilde{P}$ be any path such that (i) for all t, $W(\widetilde{P}_t) \geq \widehat{w}$ and (ii) if $W(\widetilde{P}_t) = \widehat{w}$, then $Y(\widetilde{P}_t) \geq \widehat{y}$. Then there is a strategy σ that implements $\widetilde{P}$ and is revision-proof and there is no alternative σ' that delivers the same payoffs to all players as σ and satisfies $Y(\Phi(\sigma')) > Y(\widetilde{P})$.*

Proof. Given the strengthening of (ii) to hold at $t = 1$ as well as later dates, the strategy constructed in the proof of Proposition 5 has the desired properties. $\square$

D Proof of Theorem 2

We begin with some preliminary lemmas. Define $\gamma : (\underline{y}, \overline{y}] \to (\underline{y}, \overline{y}]$ according to:

$$\gamma(y) = \begin{cases} y' \in \underset{(\underline{y},\overline{y}]}{\operatorname{argmax}}\ R^{-1}(\underline{p}, \cdot)(y) & y \in (\underline{y}, R(\underline{p}, \overline{y})] \\ y' = \overline{y} & y \in (R(\underline{p}, \overline{y}), \overline{y}]. \end{cases}$$

Thus, if $y \in (\underline{y}, R(\underline{p}, \overline{y})]$, then $y = R(\underline{p}, \gamma(y))$ and $\gamma(y)$ is next period's best continuation payoff when y is today's continuation payoff and $\underline{p}$ is played in the next period. Note that if $\beta = 0$ (e.g. in two period-lived OLG models), then R is constant in its second argument, $\underline{y} = R(\underline{p}, \overline{y})$ and $(\underline{y}, R(\underline{p}, \overline{y})]$ is empty. Thus, if $y \in (\underline{y}, R(\underline{p}, \overline{y})]$, then, necessarily, $\beta > 0$.

Lemma 5. *Let Assumption 2 hold. γ is is strictly increasing on $(\underline{y}, R(\underline{p}, \overline{y})]$. For $y \in (\underline{y}, R(\underline{p}, \overline{y})]$, $\gamma(y) \geq y + \frac{(1-\beta)}{\beta}(y - \underline{y})$.*

Proof. The strict monotonicity of γ is immediate from the definition and, by Assumption 2, the monotonicity of $R(\underline{p}, \cdot)$. For $y \in (\underline{y}, R(\underline{p}, \overline{y})]$, we have: $R(\underline{p}, \underline{y}) \leq \underline{y} < y = R(\underline{p}, \gamma(y))$. Thus, by the Lipschitz property of R in Assumption 2, $y - \underline{y} \leq R(\underline{p}, \gamma(y)) - R(\underline{p}, \underline{y}) \leq \beta(\gamma(y) - \underline{y})$ and so $\gamma(y) \geq y + \frac{(1-\beta)}{\beta}(y - \underline{y})$. $\square$

Lemma 6 constructs a function $\chi : (\underline{y}, R(\underline{p}, \overline{y})] \to (\underline{y}, \overline{y}]$ that is used to define punishment continuation payoffs in the proof of Theorem 2.

Lemma 6. *Let Assumptions 2 and 6 hold. If $R(\underline{p}, \overline{y}) > \underline{y}$, then there is a function $\chi : (\underline{y}, R(\underline{p}, \overline{y})] \to (\underline{y}, \overline{y}]$ such that i) for each $y \in (\underline{y}, R(\underline{p}, \overline{y})]$, $\chi(y) > y$ and $Q(u(\underline{p}), \gamma(y)) > Q(u(\widehat{p}), \chi(y))$, ii) if $y, y' \in (y, R(\underline{p}, \overline{y})]$ and $y' > y$, then $\chi(y') - y' \geq \chi(y) - y$.*

Proof. Let $y \in (\underline{y}, R(\underline{p}, \overline{y})]$. We first show that $Q(u(\underline{p}), \gamma(y)) > Q(u(\widehat{p}), y)$. Suppose not, i.e. $Q(u(\underline{p}), \gamma(y)) \leq Q(u(\widehat{p}), y)$. Recall that $Q(u(\underline{p}), y'(\underline{p})) = \underline{w} = Q(u(\widehat{p}), \underline{y})$. Also,

$R(\underline{p}, y'(\underline{p})) = \underline{y} < y = R(\underline{p}, \gamma(y))$, and since $R(\underline{p}, \cdot)$ is monotone by Assumption 2, $y'(\underline{p}) < \gamma(y)$. Then by the Lipschitz properties of R and Q in Assumptions 2 and 6, $Q(u(\underline{p}), \gamma(y)) - \underline{w} \leq Q(u(\widehat{p}), y) - Q(u(\widehat{p}), \underline{y}) < \kappa(y - \underline{y}) = \kappa(R(\underline{p}, \gamma(y)) - R(\underline{p}, y'(\underline{p})) \leq \kappa\beta(\gamma(y) - y'(\underline{p})) \leq Q(u(\underline{p}), \gamma(y)) - Q(u(\underline{p}), y'(\underline{p})) = Q(u(\underline{p}), \gamma(y)) - \underline{w}$. This is a contradiction. Hence, $Q(u(\underline{p}), \gamma(y)) > Q(u(\widehat{p}), y)$.

Define $d : (\underline{y}, R(\underline{p}, \overline{y})] \to (\underline{y}, \overline{y}]$ according to, for $y \in (\underline{y}, R(\underline{p}, \overline{y})]$, $Q(u(\widehat{p}), d(y)) = Q(u(\underline{p}), \gamma(y))$. By the continuity of Q in Assumption 2 and the previous inequality, $Q(u(\underline{p}), \gamma(y)) > Q(u(\widehat{p}), y)$, d is well defined and $d(y) > y$. Let $y < y' \leq R(\underline{p}, \overline{y})$. We show that $d(y) - y \leq d(y') - y'$. Assume to the contrary that $d(y') - d(y) < y' - y$. We have:

$$\begin{aligned}
0 &< Q(u(\underline{p}), \gamma(y')) - Q(u(\underline{p}), \gamma(y)) = Q(u(\widehat{p}), d(y')) - Q(u(\widehat{p}), d(y)) \leq \kappa(d(y') - d(y)) \\
&< \kappa(y' - y) \leq \kappa(R(\underline{p}, \gamma(y')) - R(\underline{p}, \gamma(y))) \leq \beta\kappa(\gamma(y') - \gamma(y)) \\
&\leq Q(u(\underline{p}), \gamma(y')) - Q(u(\underline{p}), \gamma(y)).
\end{aligned}$$

But this is a contradiction, so that $d(y) - y \leq d(y') - y'$. Set $\chi(y) = (d(y) + y)/2$. Then, since $d(y) > y$, $d(y) > \chi(y) = (d(y) + y)/2 > y$. Also, by the monotonicity of Q, $Q(u(\underline{p}), \gamma(y)) = Q(u(\widehat{p}), d(y)) > Q(u(\widehat{p}), \chi(y))$. This proves (i) in the Lemma. If $y, y' \in (y, R(\underline{p}, \overline{y})]$ and $y' > y$, then $\chi(y') - y' = (d(y') - y')/2 \geq (d(y) - y)/2 = (\chi(y) - y)$. This proves (ii) in the Lemma. □

Proof of Theorem 2.

Proof. Construct a strategy σ as follows. Set $\Phi(\sigma) = \widetilde{P}$. Proceed through successive dates $t \in \mathscr{N}$. At each date t, the continuation strategy following a defection from $\widetilde{P}$, $(\sigma|\widetilde{P}^{t-1}, p)$, $p \neq \widetilde{p}_t$, is constructed so as to incorporate a phase in which player payoffs are "built up" until they exceed $\widehat{w}$. Thereafter a revision-proof continuation strategy constructed as in the proof of Corollary 2 is played. At t, $W(\widetilde{P}_t) > \underline{w} = Q(u(\widehat{p}), \underline{y})$. By Assumption 2 and the continuity and monotonicity of Q, there is a $y_1 \in (\underline{y}, \overline{y}]$ such that $W(\widetilde{P}_t) > Q(u(\widehat{p}), y_1)$. We use a function $v : \cup_{r \in \mathscr{N}} \mathscr{P}^r \to [\underline{y}, \overline{y}]$ to attach continuation payoffs to histories during the build up phase. v is defined recursively as follows. For each $p \in \mathscr{P} \setminus \{\widetilde{p}_t\}$, let $v(p) = y_1$. If $v(P^r) \in (\underline{y}, R(\underline{p}, \overline{y})]$ and v has not previously taken values outside of $(\underline{y}, R(\underline{p}, \overline{y})]$, i.e. for all P^s such that P^r is a successor to P^s, $v(P^s) \in (\underline{y}, R(\underline{p}, \overline{y})]$, then $v(P^r, p)$ is updated according to:

$$v(P^r, p) = \begin{cases} \gamma(v(P^r)) & \text{if } p = \underline{p} \\ \chi(v(P^r)) & \text{otherwise.} \end{cases}$$

Also, $\sigma_{t+r}(\widetilde{P}^{t-1}, P^r)$ is set to $\underline{p}$. If $v(P^r)$ enters $(R(\underline{p}, \overline{y}), \overline{y}]$ for the first time, then the build up phase is concluded and for successor histories P^s to P^r, $v(P^s)$ is set arbitrarily. The continuation strategy following such a history, $(\sigma|\widetilde{P}^{t-1}, P^r)$, is set as follows. First, $\Phi(\sigma|\widetilde{P}^{t-1}, P^r)$ is set equal to $(p', \overline{p}, \overline{p}, \ldots)$, where p' lies in the interval between $\underline{p}$ and $\overline{p}$ and satisfies $v(P^r) = R(p', \overline{y})$. Such a p' exists by the convexity of $\mathscr{P}$, the continuity of R (Assumption 2) and the fact that $v(P^r) \in (R(\underline{p}, \overline{y}), R(\overline{p}, \overline{y})]$, since $\overline{y} = R(\overline{p}, \overline{y})$. Now, by the concavity of u and Q in Assumption 2 and Assumptions 4 and 5, $W(\Phi(\sigma|\widetilde{P}^{t-1}, P^r)) = Q(u(p'), \overline{y}) \geq \lambda Q(u(\underline{p}), \overline{y}) + (1-\lambda)Q(u(\overline{p}), \overline{y}) \geq \widehat{w}$ with equality only if $\lambda = 0$. Also, by Assumption 4, at each date $s \in \mathscr{N}$ and each sub-history $P^s = (p', \overline{p}, \overline{p}, \ldots, \overline{p})$, $W(\Phi(\sigma|\widetilde{P}^{t-1}, P^r, P^s)) = Q(u(\overline{p}), \overline{y}) \geq \widehat{w}$. Thus, along the continuation path $\Phi(\sigma|\widetilde{P}^{t-1}, P^r) = (p', \overline{p}, \overline{p}, \ldots)$ each player receives a payoff weakly more than $\widehat{w}$. In addition, if at any date along the path a player receives $\widehat{w}$, then this player and all her successors are playing $\overline{p}$ and the continuation payoff for this player's predecessors from this point onwards is at its maximal value $\overline{y}$. Thus, by Corollary 2, there is a revision-proof strategy σ' with $\Phi(\sigma') = (p', \overline{p}, \overline{p}, \ldots)$ and there is no alternative continuation strategy that weakly raises the payoff of all players after $(\widetilde{P}^{t-1}, P^r)$ and strictly raises the continuation payoff of the player at $(\widetilde{P}^{t-1}, P^{r-1})$. Set $(\sigma|\widetilde{P}^{t-1}, P^r) = \sigma'$.

If $\beta = 0$, then $y_1 > \underline{y} = R(\underline{p}, \overline{y})$ and so after all defections $p \neq \widetilde{p}_t$, this procedure sets $(\sigma|\widetilde{P}^{t-1}, p)$ to a common revision-proof continuation strategy σ'. In this case, define $\overline{r} = 1$. If $\beta > 0$, then let $\varepsilon_1 = \frac{1-\beta}{\beta}(y_1 - \underline{y}) > 0$ and $\varepsilon_2 = \chi(y_1) - y_1 > 0$. Now, by Lemmas 5 and 6, $v(P^{r+1}) - v(P^r) \geq \min(\gamma(v(P^r)) - v(P^r), \chi(v(P^r)) - v(P^r)) \geq \varepsilon = \min(\varepsilon_1, \varepsilon_2) > 0$. Thus, continuation payoffs rise by at least ε during each period of the build up phase and this phase cannot last more than $\overline{r} = (R(\underline{p}, \overline{y}) - y_1)/\varepsilon + 1$ periods. Thus, after all histories $(\widetilde{P}^{t-1}, P^r)$ with the first element of P^r not equal to $\widetilde{p}_t$ and $r \geq \overline{r}$, the continuation strategy $(\sigma|\widetilde{P}^{t-1}, P^r)$ is revision-proof.

We now verify that σ is revision-proof. Let P^t be a history and P'' a candidate revision path beginning at $t+1$. We consider two cases. In the first case, P^t enters a sub-game in which a revision-proof continuation strategy constructed as in the proof of Corollary 2 is played, i.e. P^t incorporates a first deviation from $\widetilde{P}$ at some date $t_0 < t$, passes through a phase in which player continuation payoffs are built up until at some date t_1 with $t_0 < t_1 \leq t$, $v(P^{t_1}) \in (R(\underline{p}, \overline{y}), \overline{y}]$. From P^{t_1} a revision-proof continuation strategy is prescribed by σ. Since P'' revises this continuation strategy, by the definition of revision-proofness, it cannot make all players from $t+1$ onwards weakly better off and some strictly better off relative to the strategy.

In the second case, P^t does not enter a sub-game in which a revision-proof continuation strategy constructed as in Corollary 2 is played, i.e. $P^t = \widetilde{P}^t$ or P^t incorporates a first deviation from $\widetilde{P}$ at some date $t_0 \leq t$ and passes through a payoff build up phase that is not concluded by t, $v(P^t) \in (\underline{y}, R(\underline{p}, \overline{y})]$. There are two sub-cases.

In the first, P'' reverts to the play prescribed by the strategy before a revision-proof continuation strategy is reached, i.e. at a date $t_1 + 1 \in \mathscr{N}$ such that $v(P^t, P''^{t_1+1}) \in (\underline{y}, R(\underline{p}, \overline{y})]$. Then, either $t_1 = 0$ and $P'' = \Phi(\sigma|P^t)$, in which case, trivially, no player is made strictly better off by the revision or $t_1 > 0$ and $P'' \neq \Phi(\sigma|P^t)$. In the latter situation, the player at $t + t_1$ is the last deviator from the play prescribed by the strategy. The choice of the y_1 and χ terms in the construction of σ ensure that this player is strictly worse off relative to reversion to the strategy. For example, if $P^t = \widetilde{P}^t$ and $t_1 = 1$, then the $t+1$-th player is the sole defector and she obtains $W(\widetilde{P}_{t+1})$ by adhering to the strategy and $Q(u(p_1''), y_1) \leq Q(u(\widehat{p}), y_1) < W(\widetilde{P}_{t+1})$ if she defects. If $P^t \neq \widetilde{P}^t$ or $t_1 > 1$, then the last defector deviates from a prescribed play of $\underline{p}$ at $t + t_1$. She forfeits a payoff of $Q(u(\underline{p}), \gamma(v(P''^{t_1-1})))$ and obtains $Q(u(p_{t_1}''), v(P''^{t_1})) = Q(u(p_{t_1}''), \chi(v(P''^{t_1-1}))) \leq Q(u(\widehat{p}), \chi(v(P''^{t_1-1}))) < Q(u(\underline{p}), \gamma(v(P''^{t_1-1})))$.

In the second sub-case, P'' does not revert to the play prescribed by σ before a revision-proof continuation strategy is reached. However, by the argument above a date $t_2 \leq \bar{r}$ is reached at which $(\sigma | P^t, P''^{t_2-1})$ is revision-proof. Then either some player from $t + t_2$ onwards is made worse off by P_{t_2}'' relative to reversion to the strategy or, by Corollary 2, all players from $t + t_2$ onwards receive the same payoff from P_{t_2}'' as from reversion to the strategy and the continuation payoff $Y(P_{t_2}'')$ is no more than that under the strategy. In the latter situation, if there are deviating players between $t + 1$ and $t + t_2 - 1$, then the last such deviator (at $t + t_1$ with $1 \leq t_1 \leq t_2 - 1$) is made strictly worse off relative to the strategy since she receives a continuation payoff $Y(P_{t_1+1}'')$ weakly less than that available under the strategy, $v(P''^{t_1})$, and so

$$
\begin{aligned}
Q(u(p_{t_1}''), Y(P_{t_1+1}'')) &\leq Q(u(p_{t_1}''), v(P''^{t_1})) = Q(u(p_{t_1}''), \chi(v(P''^{t_1-1}))) \\
&\leq Q(u(\widehat{p}), \chi(v(P''^{t_1-1}))) < Q(u(\underline{p}), \gamma(v(P''^{t_1-1}))) = W(\Phi(\sigma | P^t, P''^{t_1-1})).
\end{aligned}
$$

Thus, either some player is made worse off or no players are made better off by the revision.

These cases exhaust all possibilities. There is no history P^t and path P'' such that the latter strictly raises the payoffs of some players and does not reduce the payoff of any relative to the strategy. σ is revision-proof. □

E Sub-game perfection with quasi-recursive payoffs and state variables

We impose the following assumption on quasi-recursive games with state variables.

Assumption 2′. *(i) $\mathscr{P}$ is a closed, convex subset of a normed space $\mathscr{P}_0$ and $\mathscr{K}$ a closed, convex subset of $\mathbb{R}^m$. Endow Graph $\Pi \subseteq \mathscr{K} \times \mathscr{P}^\infty$ with the (relative) product topology.*

(ii) Λ is continuous and concave and $\Gamma : \mathscr{K} \to 2^{\mathscr{P}} \setminus \emptyset$ is compact-valued and continuous.

(iii) $u : \mathscr{K} \times \mathscr{P} \to \overline{\mathbb{R}}$ is real valued and continuous on (the closed set) Graph Γ and is otherwise $-\infty$. $Q : \overline{\mathbb{R}} \times \overline{\mathbb{R}} \to \overline{\mathbb{R}}$ is real-valued if both of its arguments are real-valued and is $-\infty$ valued otherwise. Q is increasing in both of its arguments and continuous and concave on $\mathbb{R} \times \mathbb{R}$.

(iv) $R : \mathscr{K} \times \mathscr{P} \times \overline{\mathbb{R}} \to \overline{\mathbb{R}}$ is real-valued if $(k, p, y) \in$ Graph $\Gamma \times \mathbb{R}$ and is $-\infty$ valued otherwise. For each (k, p), $R(k, p, \cdot)$ is non-decreasing. R is continuous on Graph $\Gamma \times \mathbb{R}$. There is a $\beta \in (0, 1)$, $\alpha \in (0, \beta^{-1})$ and a continuous function $\psi : \mathscr{K} \to \mathbb{R}_{++}$ such that:

(a) for all $(k,p) \in$ *Graph* Γ *and* $y, y' \in \mathbb{R}$

$$|R(k,p,y) - R(k,p,y')| \leq \beta |y - y'|;$$

(b)

$$\sup_{(k,p)\in Graph\, \Gamma} \left| \frac{\psi(\Lambda(k,p))}{\psi(k)} \right| < \alpha;$$

(c)

$$\sup_{(k,p)\in Graph\, \Gamma} \left| \frac{R(k,p,0)}{\psi(k)} \right| < \infty.$$

Given ψ as in Assumption 2′, let

$$\mathscr{C}_\psi = \left\{ f : \text{ Graph } \Pi \to \mathbb{R} \,\middle|\, f \text{ continuous and } \sup_{\text{Graph } \Pi} \left| \frac{f(k,P)}{\psi(k)} \right| < \infty \right\}.$$

The following theorem is a variant of the continuous existence theorem of Becker and Boyd (1997). It ensures the existence of a continuation payoff function Y whose restriction to Graph Π belongs to C_ψ (and is, hence, continuous and ψ-bounded.)

Theorem 3. *Let* $\mathscr{K}$, $\mathscr{P}$, Π, R *and* Λ *satisfy Assumption 2′. There exists a unique function* $Y : \mathscr{K} \times \mathscr{P}^\infty \to \mathbb{R} \cup \{-\infty\}$ *such that (a) the restriction of* Y *to Graph* Π *belongs to* $\mathscr{C}_\psi$*, (b) for* $(k,P) \notin$ *Graph* Π, $Y(k,P) = -\infty$ *and (c)* Y *satisfies for all* $(k,p,P') \in \mathscr{K} \times \mathscr{P}^\infty$,

$$Y(k,p,P') = R(k,p,Y(\Lambda(k,p),P')).$$

Proof. Available upon request. □

For $\mathscr{Y}' : \mathscr{K} \to 2^{\mathbb{R}} \backslash \emptyset$, the $\mathscr{B}$-operator is given by:

$$\mathscr{B}(\mathscr{Y}')(k) = \Bigg\{ y \in \mathbb{R} \,\Bigg|\, \exists (p,y') \in \Gamma(k) \times \mathbb{R} \text{ with } y = R(k,p,y'), y' \in \mathscr{Y}'(\Lambda(k,p)) \text{ and}$$
$$Q(u(k,p),y') \geq \sup_{p'' \in \Gamma(k)} \inf_{y'' \in \mathscr{Y}'(\Lambda(k,p''))} Q(u(k,p''),y'') \Bigg\}$$

The following analogue of Proposition 11 may then be obtained by applying similar arguments.

Proposition 12. *(i) For all* $k \in \mathscr{K}$, $\mathscr{Y}(k) = \mathscr{B}(\mathscr{Y})(k)$*. (ii) If for all* $k \in \mathscr{K}$, $\mathscr{B}(\mathscr{Y}')(k) \subseteq \mathscr{Y}'(k)$*, then for all* $k \in \mathscr{K}$, $\mathscr{Y}'(k) \subseteq \mathscr{Y}(k)$*. (iii)* $\mathscr{Y}$ *is compact-valued.*

Proof. Available upon request. □

References

Abreu, D., D. Pearce, and E. Stacchetti (1990). Towards a theory of discounted games with imperfect information. *Econometrica 58*, 1041–1063.

Ales, L. and C. Sleet (2011). Finite revision-proofness. Working Paper.

Asheim, G. (1997). Individual and collective time-consistency. *Review of Economic Studies 64*, 427–443.

Becker, R. and J. Boyd (1997). *Capital theory, Equilibrium Analysis and Recursive Utility*. Blackwell Press.

Bernheim, B. D., B. Peleg, and M. Whinston (1987). Coalition-proof nash equilibria i: Concepts. *Journal Economic Theory 42*(1), 1–12.

Bernheim, B. D., D. Ray, and S. Yeltekin (1999). Self-control, saving, and the low asset trap. Working Paper.

Caplin, A. and J. Leahy (2006). The recursive approach to time inconsistency. *Review of Economic Studies 131*, 134–156.

Farrell, J. and E. Maskin (1989). Renegotiation in repeated games. *Games and Economic Behavior 4*, 327–360.

Hammond, P. (1975). Charity: Altruism or cooperative egoism. In E. Phelps (Ed.), *Altruism, Morality and Economic Theory*. Russell Sage Foundation Publications.

Kocherlakota, N. (1996). Reconsideration-proofness: A refinement for infinite time horizon time inconsistency. *Review of Economic Studies 15*, 595–609.

Leininger, W. (1986). The existence of markov perfect equilibria in a model of growth with limited altruism between generations. *Review of Economic Studies 53*, 349–367.

Pearce, D. (1987). Renegotiation proof equilibria: Collective rationality and intertemporal cooperation. Cowles Foundation DP No. 855.

Strotz, R. (1955-1956). Myopia and inconsistency in dynamic utility maxmization. *Review of Economic Studies 23*(3), 1650180.

van Damme, E. (1989). Renegotiation-proof equilibria in repeated prisoner's dilemma. *Journal of Economic Theory 51*, 206–217.