Mitri Kitti

Subgame Perfect Equilibria in Continuous-Time Repeated Games

Aboa Centre for Economics

Discussion paper No. 120

Turku 2018

The Aboa Centre for Economics is a joint initiative of the economics departments of the University of Turku and Åbo Akademi University.

ISSN 1796-3133

Printed in Uniprint
Turku 2018

Mitri Kitti

Subgame Perfect Equilibria in Continuous-Time Repeated Games

Aboa Centre for Economics
Discussion paper No. 120
July 2018

ABSTRACT

This paper considers subgame perfect equilibria of continuous-time repeated games with perfect monitoring when immediate reactions to deviations are allowed. The set of subgame perfect equilibrium payoffs is shown to be a fixed-point of a set-valued operator introduced in the paper. For a large class of discrete time games the closure of this set corresponds to the limit payoffs of when the discount factors converge to one. It is shown that in the continuous-time setup pure strategies are sufficient for obtaining all equilibrium payoffs supported by the players' minimax values. Moreover, the equilibrium payoff set is convex and satisfies monotone comparative statics when the ratios of players' discount rates increase.

JEL Classification: C72, C73

Keywords: repeated game, continuous time, subgame perfection, equilibrium payoff set

Contact information

Mitri Kitti
Department of Economics
University of Turku
FI-20014, Finland
Email: mitri.kitti (at) utu.fi

1. Introduction

A central theme in the literature on repeated games is sustaining cooperation in long run relationships. The folk theorem asserts that all feasible and individually rational payoffs can be obtained in equilibria when the players become more patient (Fudenberg and Maskin, 1986), and more recent works characterize the set of limit payoffs when players are allowed to have differential time preferences (Lehrer and Pauzner, 1999, Chen and Takahashi, 2012, Sugaya, 2015). Players becoming more patient can be interpreted as the delay between observations of past play and reactions to vanish. In view of this interpretation, we would expect the limit set of equilibrium payoffs to correspond to the equilibrium payoffs of a continuous-time repeated game where reactions to past behavior are immediate. However, to show that this is indeed the case, a characterization of the equilibrium payoff set is necessary for continuous-time repeated games, which is the main contribution of this work.

It is well-known that that strategies defined as mappings from histories of past play into actions do not necessarily lead to well-defined expected payoffs in continuous-time framework (Anderson, 1984, Simon and Stinchcombe, 1989, Stinchcombe, 1992, Alós-Ferrer and Kern, 2015). One approach to tackle this issue is to introduce either exogenenous or endogeneous lags into the model (Bergin, 1992, Bergin and MacLeod, 1992) such that instantaneous reactions to deviations are precluded. Another way to obtain well-defined outcomes is to assume that the strategies are conditioned on a state variable (Friedman, 1994) rather than histories of past actions. The third resolution is to take the limit of strategies in discrete time (Davidson and Harris, 1981, Fudenberg and Levine, 1986).

In this work instantaneous reactions are allowed but it is assumed that the number of changes of actions during any time interval are bounded. The corresponding class of strategies is referred to as switching strategies. Similar idea has been presented by Bergin (2006), who assumes that no player can switch actions twice in an instant of time. The main purpose of this work is to characterize the induced behaviour resulting from equilibria in switching strategies and the corresponding payoff set. No such results are available for continuous time repeated games in the previous literature.

It is shown that the equilibrium payoff set is a fixed-point of a set-valued operator introduced in this work. This result parallels the classical characterization of equilibrium payoffs in discrete time repeated games (Abreu et al., 1986, 1990, Cronshaw and Luenberger, 1994) and its generalizations for stochastic games (Kitti, 2013, 2016b). An important difference to discrete time repeated games is that the equilibrium payoff set is convex, while for discrete time repeated games the payoff sets can be highly complex—even fractals (Berg and Kitti, 2013, 2016, 2014). Convexity also entails that pure strategies are sufficient for obtaining all equilibrium payoffs in continuous-time repeated games.

In this paper the players are allowed to have unequal discount time preferences, i.e, unequal discount rates. The limit payoffs of corresponding discrete time repeated games when players become more patient have raised attention in the recent literature, see Chen and Takahashi (2012) and Sugaya (2015) for extensions of the early work on the topic by Lehrer and Pauzner (1999). In essence, these papers present folk theorema for games with unequal discount factors. As shown by Kitti (2016a) the set of payoffs characterized in this paper is the limit of a class of discounted repeated games when the discount rates are constant and the discount factors approach to one.

The paper is structured as follows. Section 2 goes through the main concepts used in the paper. The characterization of the equilibrium payoff set is presented in Section 3. The action profiles and the properties of the equilibrium payoff set are analyzed in Section 4. Finally, conclusions are discussed in Section 5.

2. Switching Strategies

There are n players, $N = \{1, \dots, n\}$ denotes the set of players. The set of actions available for player $i \in N$ is A_i. Each player is assumed to have finitely many actions. The set of action profiles is $A = \times_i A_i$. As usual, a_{-i} stands for the action profile of other players than player i, and the corresponding set of action profiles is $A_{-i} = \times_{j \neq i} A_j$. Function $u : A \mapsto \mathbb{R}^n$ gives the vector of flow payoffs that the players receive when a given action profile a is played; if $a \in A$ is played, player i receives the flow payoff $u_i(a)$.

The game is played in continuous times and the players discount the future payoffs with discount rates r_i, $i \in N$. In a continuous time setup, a history of length t contains the information on all action profiles played up to time instant t. Let H^t denote the set of length t histories. A strategy profile maps histories into actions. The path of actions that a strategy profile prescribes after a given history h is called an induced path after h.

It is well known that the notion of strategy is ambiguous for continuous time games (Simon and Stinchcombe, 1989) unless the strategies are restricted, for example by introducing time lags (Bergin, 1992, Bergin and MacLeod, 1992). Hence, the strategies considered in this work have a simple structure which guarantees that histories are well-defined, and the strategies lead to unique outcomes. In particular, for a given path of play $p(t;\sigma) \in A$, for all $t \geq 0$, corresponding to a strategy profile σ, the payoffs

$$U_i(\sigma) = r_i \int_0^\infty e^{-r_i t} u_i(p(t;\sigma))dt, \ i \in N, \tag{1}$$

are well-defined.

It is assumed that a strategy profile σ has the property that the path of play after any history $h \in H^t$ is composed of a sequence (a^0, t^0), (a^1, t^1), $(a^2, t^2), \ldots$, where $t^{j-1} \leq t^j$, $j \geq 1$. The difference $t^j - t^{j-1}$ indicates how long a^j is played. Following Stinchcombe (1992), the time instants t^j, $j \geq 0$, are called jump times.

The players are allowed to switch to different actions several times during one time instant. However, it is assumed that there can be only finitely many such switches in any time interval.

Definition 1. A path of play $p(t)$, $t \geq 0$, is a switching path if there are $a^j \in A$, and t^j, $j = 0, 1, 2, \ldots$ such that $t^j \geq t^{j-1} \geq 0$ for all $j = 1, 2, \ldots$, and

1. $p(t^j) = a^j$, $a^j \neq a^{j+1}$, for all $j = 0, 1, 2, \ldots$,
2. if $t^{j+1} > t^j$ for $j = 0, 1, 2, \ldots$, then $p(t) = a^j$ for all $t \in [t^j, t^{j+1})$,
3. for any time interval $[t_0, t_1]$, $0 \leq t_0 < t_1$ there are finitely many indices j such that $t^j \in [t_0, t_1]$.

The first two assumptions state that t^j indicate jump times, and between the jump times the actions stay the same. The last condition means that there can be only finitely many switches in any time interval. In particular, for any t^j there can be only finitely many $k = 0, 1, \ldots$ such that $t^j = t^k$. This condition assures that the switching path leads to a sequence $(a^0, t^0), (a^1, t^1), (a^2, t^2), \ldots$ such that t^j goes to infinity as j increases. Hence, the payoffs corresponding to switching paths are well defined by the formula (1).

It is assumed that the players use only strategies in which they follow switching paths after all histories of past play. A history up to time instant t, denoted by h_t, is a path of past play $p(\tau)$, $\tau \in [0, t]$. The set of histories up to time instant t is H^t. It is assumed that all deviations from the ongoing path are observed immediately. For $t^j > t^{j-1}$, $j \geq 1$, and $t \in [t^{j-1}, t^j)$, a deviation means that the player takes some other action than what is prescribed by the strategy, i.e., the ongoing switching path. At jump-time t^j, $j \geq 0$, a deviation means that a player takes an action not prescribed by the corresponding action-profile a^j of the strategy. It is assumed that players are able to make only finitely many deviations at each time instant. If not, there would be ambiguity in how the game proceeds at a time instant in which some of the players deviate infinitely many times. Deviations are observed immediately, which means that the game may switch to a punishment path at the same time instant when some of the players deviate.

Definition 2. A strategy profile is a switching strategy if it induces a switching path after any history of past play.

The set of strategy profiles corresponding to switching strategies is denoted by Σ. The projection of switching strategies to the strategies of player i is Σ_i. For $\sigma \in \Sigma$, the individual strategy $\sigma_i \in \Sigma_i$ stands for the strategy of player i given that the other players choose their actions according to σ. The restriction of a switching strategy after history $h_t \in H^t$, $t \geq 0$, is $\sigma|h_t$. By definition, a switching strategy induces a path $p(s)$, $s \geq t$, which consists of action profile-jump time pairs $(a^0, t^0), (a^1, t^1), (a^2, t^2), \ldots$ after any history h_t. Note that strategies in discrete-time repeated games can be embedded into switching strategies by requiring that the players' actions stay the same during a given time interval $\Delta > 0$ and they can be changed only in the discrete time instants determined by the time step Δ.

Example 1. To clarify the concept of induced path let us consider the games with payoffs parameterized by a scalar α defined as below.

	L	R
T	α, α	$-1, 1$
B	$1, -1$	$-2, -2$

Assume that the strategy σ_1 of player 1 (the row player) is to start by playing B for time interval Δ. After time period of length Δ that has passed from player's previous switch, the player 1 changes to T if player 2 has played L during the previous time interval of length Δ. If player 2 has played R during the previous time interval while player 1 has played B, player 1 chooses B for the next time period of length Δ. The strategy σ_2 of player 2 is defined analogously but assuming that player 2 starts by choosing L and then switches between R and L.

The strategy profile σ described above induces infinitely many paths. First, if no-player deviates the path $p(t)$ is the one in which $a = (B, L)$ is played in every second interval of length Δ beginning from $t = 0$, and $b = (T, R)$ is played otherwise. To be exact, a is played during the time intervals $[0, \Delta)$, $[2\Delta, 3\Delta)$, $[4\Delta, 5\Delta)$, ... On the other hand, history in which no player deviates induces also paths that correspond to restrictions of $p(t)$ that begin from any other time instant than $t = 0$. Histories in which players deviate induce similar paths. All the induced path in this example are switching paths.

One way to define strategies that lead to switching paths is given by Bergin (2006), who proposes two conditions for strategies in continuous time games. First, players' strategies should be right constant if the history is not left continuous, and second, the strategies should be constant after a point of time at which the history is left continuous. These strategies belong to the class of switching strategies. However, the class of switching strategies also contains strategies in which multiple actions can be taken at a single time instant, which is not allowed in the framework of Bergin (2006). Allowing for taking multiple actions at the same time instant is not completely new idea. Murto and Välimäki (2013) define a continuous-time game in which players can react to each others' actions immediately and exit the game after observing some of the players exiting at the same time instant. As in this game, the switches of actions divide the game into separate stages during which the actions stay constant.

What is common for switching strategies and the strategies introduced by Bergin (2006) is that there is an identifiable first time for any change of action. This is a central feature that is required when reactions to deviations are allowed to be immediate. For example, assume that the players follow a path in which action profile $a \in A$ is played for all $t \neq t_0$, and for $t = t_0$ the action profile is $b \in A$, $b \neq a$. If player, $i \in N$ deviates and chooses $a_i \in A_i$ for $t \in [t_0, t_1)$, then there would not be any identifiable first time instant at which the player started the deviation.

The subgame perfection of a switching strategy-profile is defined in the usual way: after any history none of the players is willing to deviate from the play prescribed by an equilibrium strategy.

Definition 3. The strategy profile $\sigma \in \Sigma$ is a subgame perfect equilibrium if

$$U_i(\sigma|h_t) \geq U_i(\sigma_i', \sigma_{-i}|h_t) \text{ for all } h_t \in H^t,\ t \geq 0,\ \sigma_i' \in \Sigma_i, \text{ and } i \in N.$$

3. Equilibrium Payoffs

In this section the focus is on the set of equilibrium payoffs

$$V = \{U(\sigma|h_t) : \sigma \in \Sigma \text{ is subgame perfect equilibrium}, h_t \in H^t,\ t \geq 0\}.$$

The primary purpose is to relate V and its closure to the fixed-points of a particular set-valued operator.

3.1. Set-valued operator for payoffs

Before defining the set-valued operator some notations and definitions are needed. In the following, e^{-rt} denotes the matrix that has $e^{-r_i t}$, $i \in N$, on its diagonal. If $a \in A$ is played from time instant 0 to t after which the players get continuation values corresponding to a vector $v \in \mathbb{R}^n$, the discounted average is

$$DU_a^t(v) = (I - e^{-rt})u(a) + e^{-rt}v,$$

where I is the $n \times n$ identity matrix. Hence, the payoffs that can be obtained for the pair a and v for different values of t lie on the (forward) orbit of v under DU_a. This orbit can be viewed as a trajectory of the continuous time dynamical system, flow of which is $DU_a^t(v) = (I - e^{-rt})u(a) + e^{-rt}v$.

Definition 4. The orbit of v under DU_a is

$$\phi_{av} = \text{cl}\left\{DU_a^t(v) : t \geq 0\right\}.$$

Observe that as t goes to infinity, the point in ϕ_{av} goes to a. By taking the closure this point gets included in the orbit.

In the following v_i^- is the minimax payoff of player i in the stage game and $v_i^-(V) = \inf\{w_i : w \in V\}$ for $i \in N$. The minimax payoffs are in mixed strategies and it is assumed that the players are able to randomize their actions in for each time interval, i.e., they can choose randomly their own action for each time interval $[t_1, t_2]$, $t_2 > t_1 \geq 0$. Consequently they can guarantee themselves their minimax payoffs. As will be observed in Section 4.2, all the payoffs supported by the minimax values can be attained in pure strategies. Hence, the main attention will be on pure strategies. Note that $v_i^-(V) \geq v_i^-$ for all $i \in N$. The payoff vector corresponding to minimax payoffs is v^-. The set of payoffs in $\mathbb{R}^n$ for which $w \geq v^-$, i.e., the individually rational payoff set, is denoted by IR.

Assume players are asked how long they would accept playing an action profile $a \in A$ given that it will be followed by a continuation payoff v. The answer would be the largest $t > 0$ such that

$$(I - e^{-rt})u(a) + e^{-rt}v \geq v^-.$$

If for some player $v_i^-(V) > v_i^-$, then we could replace the right hand side v_i^- with $v_i^-(V)$ for the player. However, the only reason why it may happen that $v_i^-(V) > v_i^-$, is that there is always some other player $j \neq i$ for whom $(1 - e^{-r_j t})u_j(a) + e^{-rt}v_j = v_j^-$ before this condition would hold for player i. Hence, the relevant incentive compatibility condition for playing a and then letting the players receive the continuation value vector v is

$$DU_a^\tau(v) \geq v^- \text{ for all } \tau \in [0, t].$$

Example 2. To illustrate orbits consider a game as defined in Example 1 for $\alpha = 2$. Let us label the action profiles as follows: $a = (T, L)$, $b = (T, R)$, $c = (B, L)$, and $d = (B, R)$. Corresponding orbits are illustrated in Figure 1 for $v = u(B, R)$.

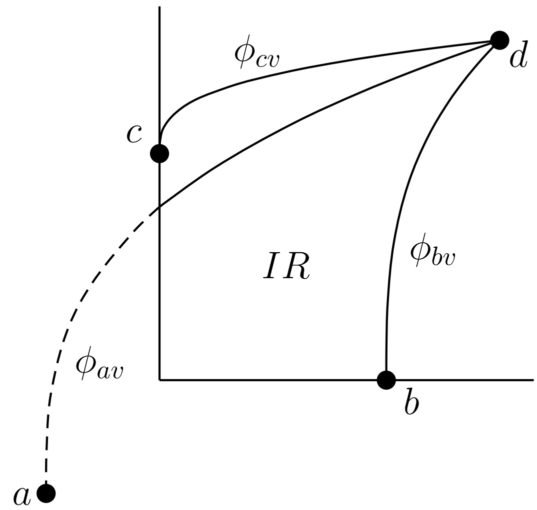

Figure 1: Examples of orbits for $r_1 = 3$, $r_2 = 1$.

The next step is to define an operator that generates all the possible payoffs that can be obtained from a compact set $W \subseteq IR$ by playing the action profiles for all possible time periods. The relevant operator takes the payoffs from the orbits ϕ_{av} that satisfy the incentive compatibility conditions $w \geq v^-$:

$$B_a(W) = \bigcup_{v \in W} \{w \in \phi_{av} : w \geq v^-\},$$

and by taking the union over A we get the set-valued operator

$$B(W) = \bigcup_{a \in A} B_a(W).$$

Note that when a is played as long as possible given that the continuation payoff after playing a is v, the resulting payoffs are in ϕ_{av}. In particular, assume for a while that $W = V$. Taking a point on $\phi_{a^1 v^1}$ means that it is specified how long a^1 is played, and after that time period the players receive the continuation payoffs corresponding to vector v^1. Let us say that t^1 indicates how long a^1 is played. Because v^1 belongs to V, it corresponds to an action profile $a^2 \in A$, a jump time t^2, and a continuation payoff v^2. Repeating this construction gives a path $(a^1, t^1), (a^2, t^2), \ldots$.

An important property of the operator B is that it is non-decreasing in the set-inclusion; $W \subseteq B(W)$ for any set $W \subset IR$. This follows by observing that W is obtained as the image of B_a, $a \in A$, because $B_a(v)$ contains v (point corresponding to $t = 0$ in the orbit ϕ_{av}).

Our first result on V is that it is a fixed-point of B.

Proposition 1. $V = B(V)$.

Proof. By the monotonicity of B, it holds that $V \subseteq B(V)$. If the inclusion would be strict, then there would be $t > 0$, $a \in A$, and $v \in V$ such that $DU_a^\tau(v) \geq v^-$ for all $\tau \in [0, t]$, and $w = DU_a^t(v) \notin V$. Hence, the switching path p defined by setting $p(\tau) = a$ for $\tau \in [0, t)$ after which $p(\tau)$ equals the path p' that gives v, would be an induced equilibrium path of a strategy constructed as follows. After t the players follow the strategy corresponding to p' and v. If during the time interval $[0, t)$ any player would deviate, the deviation leads to either p or p' depending which of the two would lead to a smaller payoffs to the deviator. Let p^i stand for the resulting punishment path for player $i \in N$.

The path p would lead to a smaller payoff for player i if $u_i(a) < v_i$ and otherwise p' would be worse for player i. No player is willing to deviate during $[0,t)$ because $v_i(p,\tau) \leq v_i(p^i,0)$ for all i and $\tau \in [0,t)$. Here we denote

$$v_i(p,t) = r\int_0^\infty u_i(p(t+\tau))e^{-r_i\tau}d\tau.$$

If $u_i(a) < v_i$, then $v_i(p,\tau)$ for $t \in (0,t)$ is better for player i than $v_i(p^i,0)$ because it entails a faster switch to a desirable payoff v_i. On the other hand, if $u_i(a) \geq v_i$, then $v_i(p,\tau)$ is preferred to $v_i(p^i,0)$ because the latter would entail an immediate switch to an undesirable payoff compared to $u_i(a)$.

Hence, p is induced by a strategy from which no player would be willing to deviate at any point of time, i.e., p is an induced equilibrium strategy. However, this would be in contradiction with w, the payoff corresponding to p, not belonging to the set of equilibrium payoffs V. Hence, the set-inclusion $V \subseteq B(V)$ cannot be strict, which means that $V = B(V)$. □

In general B has multiple fixed-points due to the monotonicity of B. Hence, a more thorough characterization of V is needed. The second observation, demonstrated in the following example, is that V may not be a closed set.

Example 3. In this example the flow payoffs $u(a)$, $a \in A$, are as below.

	L	C	R
T	$-1,-1$	$0,0$	$-1,2$
M	$0,0$	$0,0$	$0,0$
B	$2,-1$	$0,0$	$-1,-1$

Let us denote $a = (T,R)$ and $b = (B,L)$. The discount rates are $r_1 = 10$ and $r_2 = 1$. The resulting payoff set is illustrated in Figure 2. It should be observed that the Pareto frontier in the figure does not belong to V, i.e., V is not a closed set. This is because there is no payoff point in V that would map into the corner point v illustrated in the figure, i.e., $B(v) = \{v\}$ and $v \notin B(w)$ for all $w \in V$, $w \neq v$. For points in the neighborhoods of v there are, however, payoffs that map into these points under B. For example, the orbit ϕ_{aw} in the figure, passes through a neighborhood of v.

To show more formally that v does not have any pre-image in the interior of V consider the tangent directions of ϕ_{av} and ϕ_{bv} at v. First, it should be noted that $v = \lambda u(a) + (1-\lambda)u(b)$ for $\lambda \in (0,1)$ such that $v_1 = v_1^- = 0$. This follows from the comparative statics result of Corollary 1 in Section 4.1. Moreover, the Pareto frontier is determined by ϕ_{bv}. The tangent directions of ϕ_{av} and ϕ_{bv} at v are $T_a = (1-\lambda)R[u(b) - u(a)]$ and $T_b = \lambda R[u(b) - u(a)]$, respectively. Here R denotes the matrix with discount rates on the diagonal. Hence, T_a and T_b point to opposite directions. It follows that ϕ_{av} intersects ϕ_{bv} at single point v. Hence, it is impossible to reach v by any switching path.

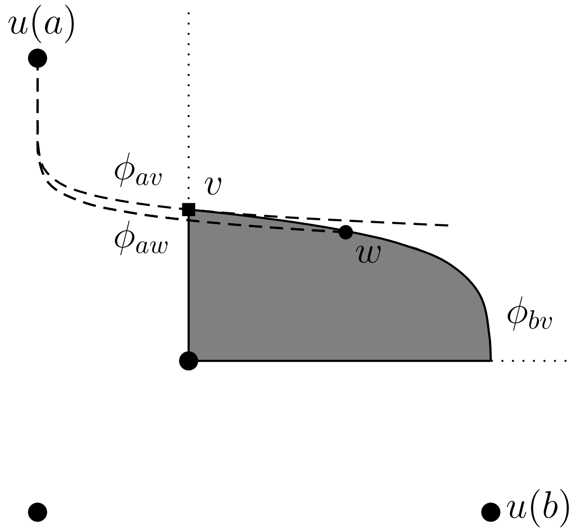

Figure 2: Example of a payoff set for $r_1 = 10$, $r_2 = 1$.

Due to possible non closedness of V, the attention in the rest of the paper is on its closure, i.e., cl(V). Importantly, this set is the limit of equilibrium payoffs of a large class of discrete time repeated games when the discount factors go to one while discount rates remain constant (Kitti, 2016a).

3.2. The closure of the equilibrium payoff set

Recall that the operator B is defined by taking the union of $B_a(W)$ over all $a \in A$. However, it may happen that not all the action profiles can be played given the continuations in the set V. To be more specific, an action profile $a \in A$ and a continuation payoff $v \in V$ can be played if there is $t > 0$ such that

$$DU_a^\tau(v) \geq v^-$$

for all $\tau \in [0, t]$. The largest set of action profiles, in terms of set inclusion, which have this property is the set of action profiles that can be played in equilibrium. Related to a set $A' \subseteq A$ we can define the space of closed sets from which the continuations can be taken. Formally this space is defined below.

Definition 5. A closed set $S \subseteq IR$ belongs to $\mathcal{C}(A')$, $A' \subseteq A$, if for any $a \in A'$ there is $v \in S$ and $t > 0$ such that $DU_a^\tau(v) \geq v^-$ for all $\tau \in [0, t]$.

The largest set of action profiles that can be played in equilibrium is the set A^* defined as follows.

Definition 6. A^* is the largest subset of A such that there is $\bar{V} \in \mathcal{C}(A^*)$ for which $W \subset B(W) \subseteq \bar{V}$ for all $W \subset \bar{V}$, $W \in \mathcal{C}(A^*)$.

Example 4. The purpose of this example is to clarify the definition of $\mathcal{C}(A')$. The payoffs are as in Example 3. The minimax payoff vector is $v^- = (0, 0)$. For the payoff set $S = \{v^-\}$ the set of action profiles that can be played for a positive period of time is $A' = \{(M, L), (M, C), (M, R), (T, C), (B, C)\}$. Each of these action profiles yields the payoff vector $(0, 0)$. Hence, $\mathcal{C}(A') = \{\emptyset, \{v^-\}\}$.

Next consider the line segment between the points $u(T, R) = (-1, 2)$ and $u(B, L) = (2, -1)$ in which the payoffs are at least $(0, 0)$, i.e., the set $W = \text{conv}\{(1, 0), (0, 1)\}$. Now the set of action profiles that can be played for a positive period of time are A (the set of all action profiles). For the set W it holds that $W \subset B(W)$. Moreover, there is a fixed-point of B containing W and this fixed-point is the appropriate $\bar{V}$ in the definition of A^*. In the specific case when the discount rates are equal, $\bar{V}$ is the convex hull of payoff vectors $(1, 0)$, $(0, 1)$, and $(0, 0)$.

The main result of this paper is that the closure of the equilibrium payoff set V is $\bar{V}$ as in the definition of A^*. Moreover, this set is the unique smallest fixed-point of the operator B in $\mathcal{C}(A^*)$. The proof of the result is presented in the Appendix.

Proposition 2. $cl(V) = \bar{V}$.

The key property that is required from $\bar{V}$ is that the condition $W \subset B(W)$ holds for any subset $W \in \mathcal{C}(A^*)$ of $\bar{V}$. In essence, this condition means that any subset of the equilibrium payoff set can be used for creating new equilibrium payoffs. In particular, it follows that this set is a fixed-point of B.

Proposition 3. *The set $\bar{V}$ is the unique smallest fixed-point of B in $\mathcal{C}(A^*)$.*

Proof. Take V^* as the intersection of all $V' \in \mathcal{C}(A^*)$ such that $V' = B(V')$. If V^* has a strict subset $W \in \mathcal{C}(A^*)$, such that $W \subset B(W)$ and $W \in \mathcal{C}(A^*)$, then $W \subset B(W) \subseteq V'$ for all fixed-point sets $V' \in \mathcal{C}(A^*)$, because otherwise there would be a contradiction with V^* being the intersection of all closed fixed-points of B in $\mathcal{C}(A^*)$. Assume that there is another set $\hat{V} \neq V^*$ with the property $W \subset B(W) \subseteq \hat{V}$ for all $W \subset \hat{V}$ and $W \in \mathcal{C}(A^*)$. Observe that $\hat{V}$ would be a fixed-point of B, which implies that V^* would be its strict subset. However, this would be in contradiction with the property $W \subset B(W) \subseteq \hat{V}$ for all $W \subset \hat{V}$, $W \in \mathcal{C}(A^*)$. Hence, $V^* = \bar{V}$ is the unique set as claimed. □

4. Properties of Equilibrium Action Profiles and Payoff Sets

4.1. Comparative statics

In repeated games unequal discount rates have the effect that payoff vectors outside the feasible and individually rational payoffs can be obtained in equilibrium (Lehrer and Pauzner, 1999). In order to understand this phenomenon we need consider the comparative statics of the set $\bar{V}$ when the discount rates are varied. In this section it is shown that the more different the players time preferences are, the larger the equilibrium payoff set.

Because the discount rates $r = (r_1, \ldots, r_n)$ are varied, the set $\bar{V}$ is denoted as a function of r, i.e., $\bar{V}(r)$ stands for the equilibrium payoffs in the game where the discount rates are determined by the vector r. The first observation on $\bar{V}(r)$ is that it only depends on the ratios of the discount rates simply because an orbit ϕ_{av} is the same whenever the ratios of the discount rates remain the same.

Proposition 4. *$\bar{V}(r)$ is homogenous of degree zero as a function of r; $V(\lambda r) = \bar{V}(r)$ for all $\lambda > 0$.*

Proof. The result follows from the observation that the tangent direction of an orbit remains the same when r is multiplied by any positive constant. The tangent vector at point t is $(r_i e^{-r_i t} u_i(a) - r_i e^{-r_i t} v_i)_{i \in N}$. Multiplying all r_i, $i \in N$ by $\lambda > 0$ only scales the tangent vector. Hence, the tangent direction of ϕ_{av} remains the same at any point on the orbit when r is multiplied, which means that the orbit itself remains the same. Hence, $B(\cdot; r)$, i.e., the fixed-point mapping for given discount rates r, remains the same when r is multiplied by a positive constant, i.e., $\bar{V}(r)$ is homogeneous of degree zero. □

As stated in the following result, the set of equilibrium payoffs increases when the ratios of the discount rates increase. This result is important because it can be used for analyzing the limit set of payoffs corresponding discrete time repeated games when the discount factors converge to one Kitti (2016a). In the following, $\pi = i_\pi^1 i_\pi^2 \cdots i_\pi^n$ denotes a permutation of players and $R(\pi)$ stands for the set of vectors of discount rates such that $r_{i_\pi^j} \geq r_{i_\pi^{j+1}} > 0$ for all $j = 1, \ldots, n$. The proof of Proposition 5 is presented in the Appendix.

Proposition 5. *If r and q are vectors of discount rates belonging to $R(\pi)$ such that $r_{i_\pi^j}/r_{i_\pi^{j+1}} \geq q_{i_\pi^j}/q_{i_\pi^{j+1}}$ for all $j = 1, \ldots, n-1$, then $\bar{V}(q) \subseteq \bar{V}(r)$.*

As a corollary of the above result, it follows that the set of feasible and individually rational payoffs of repeated games with equal discount factors belong to $\bar{V}(r)$ for all discount rates r. Let $FP = \text{conv}\{u(a) : a \in A\}$ stand for the feasible payoff vectors. Observe, in particular that $FP \cap IR$ is the set of equilibrium payoffs for continuous-time repeated games with equal discount rates.

Corollary 1. *$FP \cap IR \subseteq \bar{V}(r)$ for all vectors of discount rates r.*

4.2. Convexity of the equilibrium payoff set

The main result of this section is that the equilibrium payoff set V is convex. Convexity is an important property, because it implies that mixed or correlated strategies do not produce any additional payoffs. In essence, pure strategies are sufficient to obtain all equilibrium payoffs. The proof of the result is presented in the Appendix.

Proposition 6. *V is a convex set.*

Assume that the players use randomized strategies instead of pure strategies at each time instant, either because they have access to a public correlation device or they use mixed strategies. In this case it is assumed that the history of distributions used in the randomization is observed. In other words, players condition their actions on the history of past distributions rather than realized actions. See Berg and Schoenmakers (2017) and Berg (2018) on constructing the payoff sets corresponding to subgame-perfect mixed-strategy equilibria.

Let $\Delta(A)$ denote the distributions over A, i.e., the correlated strategies over pure actions. In case of mixed strategies we could define $\Delta(A)$ as $\Delta(A_1) \times \cdots \times \Delta(A_n)$, where $\Delta(A_i)$ is the set of mixed strategies over player i's actions. Because mixed-strategies are a subset of correlated strategies, the following operator for payoff sets is formulated only for correlated strategies. Note that the minimax payoffs in mixed strategies are

$$v_i^- = \min_{m_{-i}\in\Delta(A_{-i})} \max_{m_i\in\Delta(A_i)} \sum_{a_{-i}\in A_i} \sum_{a_i\in A_i} m_{-i}(a_{-i}) m_i(a_i) u_i(a_i, a_{-i}),$$

where $m_{-i}(a_{-i})$ denotes the probability of a_{-i} and $m_i(a_i)$ the probability of $a_i \in A_i$. When the players are able to randomize their actions for each time interval, the players can guarantee themselves their minimax payoff levels in the continuous-time repeated game.

Let us define the operator B as before except that now the union is over $B_m(W)$, where $m \in \Delta(A)$ and $m(a)$ is the probability of $a \in A$, and ϕ_{mv} in the definition of $B_m(v) = \phi_{mv} \cap IR$ is defined as

$$\begin{aligned}\phi_{mv} &= \mathrm{cl}\left\{(I - e^{-rt})\left(\sum_{a\in A} m(a)u(a)\right) + e^{-rt}v : t \geq 0\right\}\\ &= \mathrm{cl}\left\{\sum_{a\in A} m(a)\left[(I - e^{-rt})u(a) + e^{-rt}v\right] : t \geq 0\right\}.\end{aligned}$$

The following result tells that applying B_m to any payoff vector from V generates vectors in V. Hence, using randomized strategies does not create any additional payoffs, which means that pure strategies are sufficient for obtaining all payoffs from the interior of the equilibrium payoff set even when randomized strategies are allowed. All that is needed is that the players can choose their actions randomly—in practice they do not need to.

Proposition 7. *For any $m \in \Delta(A)$ and $v \in \bar{V}$ it holds that $B_m(v) \in \bar{V}$.*

Proof. By the definitions of B_m and ϕ_{mv}, it holds that

$$B_m(v) \in \mathrm{conv}\left(\cup_{a\in A}(\phi_{av} \cap IR)\right).$$

By Proposition 1, we have $\phi_{av} \cap IR \subseteq V$ for any $v \in V$. Due to the convexity of V it holds that

$$\mathrm{conv}\left(\cup_{a\in A}(\phi_{av} \cap IR)\right) \subseteq V$$

when $v \in V$. Hence, $B_m(v) \in V$. □

5. Conclusions

The strategies studied in this work allow for immediate reactions to deviations from ongoing paths of play in continuous-time repeated games. The main result is that the set of payoffs corresponding to subgame perfect equilibrium strategies can be characterized in the spirit of Abreu et al. (1990) as the smallest fixed-point of a particular set-valued mapping. This payoff set is highly relevant: it is the limit of equilibrium payoffs of a class of discrete time repeated games, when the discount factors converge to one, and the players have constant discount rates (Kitti, 2016a).

As observed in this paper, continuous-time repeated games lead to payoff sets that have a number of attractive features compared to discrete-time repeated games. For instance, the payoffs satisfy monotone comparative statics when the players discount rates become more unequal, in the sense that the ratios of discount rates increase. Moreover, the equilibrium payoff set is convex, which implies that pure strategies are sufficient for obtaining all payoffs supported by the players' minimax values.

Appendix A. Auxiliary Proofs

The following result is needed in the proof of Lemma 2, which in turn is needed to show Proposition 2.

Lemma 1. *For any initial set $W^0 \in \mathcal{C}(A^*)$ such that $W^0 \subseteq \bar{V}$ the iteration*

$$W^{k+1} = B(W^k),\ \ k = 0, 1, \ldots,$$

converges to $\bar{V}$.

Proof. Recall that B is non-decreasing; $B(W) \subseteq W$ for all $W \in \mathcal{C}(A^*)$. It follows that $W^k \subseteq W^{k+1} \subseteq \bar{V}$ for all $k = 0, 1, \ldots$. The limit set $\bar{W} = \mathrm{cl}(\cup_k \{W^k\})$ is a fixed-point of B. The limit is considered in the Painlevé-Kuratowski convergence, in which the limit is always a closed set (Rockafellar and Wets, 1998). If $\bar{W}$ was different from $\bar{V}$, we would have a contradiction with $\bar{V}$ being the smallest fixed-point of B in $\mathcal{C}(A^*)$. □

Lemma 2. *Let $W \in \mathcal{C}(A')$ for $A' \subseteq A^*$. If $S \subset B(S)$ for all $S \in \mathcal{C}(A')$ such that $S \subset W$, then $W \subseteq \bar{V}$.*

Proof. There are two alternatives: either $B(S) \subset W$ for all $S \subset W$ such that $S \in \mathcal{C}(A')$, or $W \subset B(W)$. In the former case W would be a fixed-point of B, which implies that $W \subseteq \bar{V}$ because $A' \subseteq A^*$. Hence, assume that $W \subset B(W)$.

Take any set $S \subset W^1 = B(W)$ such that $S \in \mathcal{C}(A')$. The purpose is to show that $S \subset B(S)$. First, if $S \subset W$, then by the assumption of the lemma $S \subset B(S)$. Hence, assume that S contains some point v not in W. The definitions of S, $B(W)$, and the assumption $S \subset B(W)$ imply that there are $a \in A'$, $v' \in W$, and $t > 0$ with the following properties:

1. $v = DU_a^t(v') \geq v^-$, and
2. there is $\tau > t$ such that $DU_a^\tau(v') \geq v^-$ and $DU_a^\tau(v') \notin S$.

First, note that if the latter condition fails, then for any $a \in A'$ and $v' \in W$, the whole set $\phi_{av'} \cap IR$ would belong to S, which cannot be possible because $S \subset B(W)$ and $W \subset B(W)$. Second, by the definition of DU_a^t it holds that $DU_a^\tau(v') = DU_a^{\tau-t}(v)$, which implies that $DU_a^{\tau-t}(v) \notin S$. Hence, for any subset S of $B(W)$ we have $S \subset B(S)$.

Now the same deduction made in the beginning for W can be repeated for $W^1 = B(W)$. This construction leads either to a set $W^k = B^k(W) \subseteq \bar{V}$ for some k, or to a sequence $W^k = B^k(W)$, $k = 0, 1, 2, \ldots$ with $W^0 = W$, converging to a fixed-point of B by Lemma 1. This fixed-point contains the fixed-point of B in $\mathcal{C}(A')$ because at least the actions of A' can be played for W^k for all k. Because $A' \subseteq A^*$ the fixed-point set is contained in $\bar{V}$. Hence, in either case we obtain $W \subseteq \bar{V}$. □

Proof of Proposition 2. The inequality $\bar{V} \subseteq \mathrm{cl}(V)$ follows from V being a fixed-point of B (Proposition 1), and $\bar{V}$ being the intersection of all fixed-points in $\mathcal{C}(A^*)$. Note that $\mathrm{cl}(V)$ is a fixed-point of B in $\mathcal{C}(A^*)$ when V is a fixed-point set.

It remains to be shown that $\mathrm{cl}(V) \subseteq \bar{V}$. Take any equilibrium path with the corresponding payoff vector $w(0) \in \mathcal{C}(A^*)$ in time instant $t = 0$. Let $w(t)$ stand for the equilibrium payoff as a function of time along the equilibrium path. The purpose is to show that $S \subset B(S)$ holds for any subset of $W = \mathrm{cl}\{w(t) : t \geq 0\} \in \mathcal{C}(A^*)$.

First assume that $w(t) = w$ for all $t \geq 0$. Now there is $a \in A^*$ such that $u(a) = w$, which would mean that $u(a) \geq v^-$. This would, however, imply that $u(a)$ belongs to $\bar{V}$, because such a point belongs to any fixed-point set of B. This is because for any point $w' \geq v^-$ and any $t \geq 0$ we have $DU_a^t(w') \geq v^-$.

Next assume that $w(t)$ is not constant, in which case there is a non-empty set $S \subset W$ such that $S = \{w(t) : t \in I\}$ where I is a closed subset of $\{t \in \mathbb{R} : t \geq 0\}$. Note that by definition, for any t^1 and t^2 such that $t^2 > t^1$ the payoff $w(t^1)$ is obtained from any $w(t^2)$ by taking a sequence of operations $w^{k+1} = DU_{a^k}^{\tau_k}(w^k) \geq v^-$, where $a^k \in A^*$, $w^0 = w(t^2)$, $\tau_k > 0$ and $k = 0, \dots, K$. In other words, $w(t^1)$ is contained in $B^k(\{w(t^2)\})$ for some $k \in \mathbb{N}$, which implies that $\{w(t^2)\} \subset B(\{w(t^2)\})$.

The above observation means that if there is a closed interval $I' = [t^1, t^2]$, $t^2 > 0$, that does not belong to I and $\{w(t) : t \in I'\}$ does not belong to S, then all the payoffs corresponding to the interval I' are obtained from the points in S by using the mapping B. Hence, all the sets $S \subset W$ having such missing time intervals I' satisfy $S \subset B(S)$. The remaining case is the one where I is a closed interval from 0 to T for some $T > 0$. Hence, assume that I is such a time interval.

Having $B(\{w(t)\}) \subseteq S$ for all $t \in I$, i.e., $B(S) \subseteq S$, would mean that it is not possible to obtain any other payoffs from S than the payoffs in S itself. This would, however, mean that $u(a)$ is constant for all $a \in A^*$ and $w(t)$ would be constant, too. Otherwise $DU_a^t(w(0)) \neq DU_b^t(w(0))$ for $a, b \in A^*$, such that $u(a) \neq u(b)$ and $DU_a^\tau(w(0)), DU_b^\tau(w(0)) \geq v^-$ for $\tau \in [0, t]$ for some t. Note that the existence of suitable $t > 0$ follows from $w(0) \in \mathcal{C}(A^*)$. Note also that $B(S) \subseteq S$, would mean that $S = B(S)$. Hence, the condition $S \subset B(S)$ holds for all $S \in \mathcal{C}(A^*)$ such that $S \subset W$, which by Lemma 2 implies that W belongs to $\bar{V}$.

If W would correspond to the set of payoffs along an equilibrium path such that $w(0) \in \mathcal{C}(A')$ for $A' \subset A^*$, we could deduce by the above arguments that W would belong to some fixed-point set of B contained in $\bar{V}$. Hence, any set of payoffs W corresponding to an equilibrium path is a subset of $\bar{V}$, which implies that $\mathrm{cl}(V) \subseteq \bar{V}$. □

Proof of Proposition 5. The proof is based on showing the inclusion $\bar{V}(q) \subseteq B(\bar{V}(q); r)$. In the following, the set $\phi_{av}(r)$ stands for the orbit of v under DU_a for discount rates r. In the proof it is assumed that the players' smallest equilibrium payoffs are the minimax payoffs v_i^-, $i \in N$. However, if this was not the case, then it follows from the below proof that the smallest equilibrium payoffs for discount rates r are no larger what they are for discount rates q. Hence, the result of the proposition holds also when some players' smallest equilibrium payoffs are larger than the minimax payoffs.

Let $\partial\bar{V}(q)$ stand for the relative boundary of $\bar{V}(q)$. Take any point $v^* \in$

$\partial\bar{V}(q)$ such that $v_i^* > v_i^-$, $i \in N$. There is $a \in A^*$ corresponding to v^* such that the tangent direction of $\phi_{av^*}(q)$ at v^* (i.e., at $t = 0$) does not point into the set $\bar{V}(q)$. This follows from the convexity of $\bar{V}(q)$, see Proposition 6 in Section 6, and is illustrated in Figure A.3.

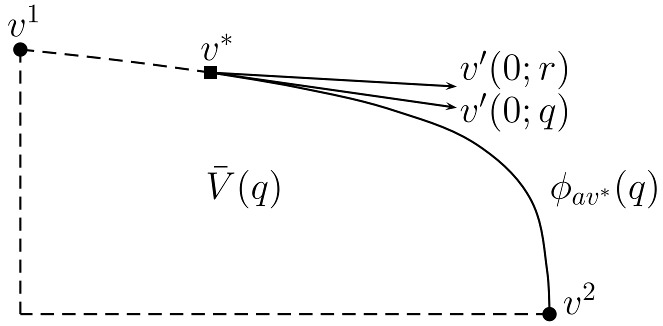

Figure A.3: Illustration of the arguments in the proof of Proposition 5 when $r_1 > r_2$ and $q_1 > q_2$.

Take the tangent direction of $\phi_{av^*}(r)$ at $t = 0$. Recall that the points along $\phi_{av^*}(r)$ can be presented as $v(t;r) = (I - e^{-rt})u(a) + e^{-rt}v^*$, $t \geq 0$. Hence, the components of the tangent vector $v'(0;r)$ at origin are $dv_i(0;r)/dt = r_i[u_i(a) - v_i^*]$, $i = 1, \ldots, n$. Without loss of generality, we may assume that $r_1 \geq r_2 \geq \ldots \geq r_n$. It follows that the absolute value of $dv_i(0;r)/dv_n(0;r) = r_i[u_i(a) - v_i^*]/[r_n(u_n(a) - v_n^*)]$, $i \neq n$, is greater for r than for q. Hence, the orbit $\phi_{av^*}(r)$ points outwards from $\bar{V}(q)$.

Let us next consider (relative) boundary points v^* such that $v_i^* = v_i^-$ for some $i \in N$, i.e., corner points of $\bar{V}(q)$. There are now two alternatives: either $v(t;q) \in \partial\bar{V}(q)$ for some $t > 0$ (and $a \in A$) or $v(0;q) \in \partial\bar{V}(q)$, i.e., $v(0;q)$ is the only point on the orbit ϕ_{av^*} that belongs to $\partial\bar{V}(q)$ for any $a \in A$. These two cases are illustrated in Figure A.3: for point $v^1 \in \partial\bar{V}(q)$ the payoff $v(t;q)$ stays on the boundary $\partial\bar{V}(q)$ for $t > 0$, whereas v^2 has the property that it is the only point on $\partial\bar{V}(q)$ on any orbit $\phi_{av^2}(q)$ for $a \in A$.

For points of type v^1 the previous deduction applies, while points of type v^2 satisfy $v^2 \in B(\bar{V}(q);r)$ because $v^2 = DU_a^0(v^2) \geq v^-$. To gather, the points of $\partial\bar{V}(q)$ either stay the same or are mapped outwards from $\partial\bar{V}(q)$ under the mapping $B(\cdot;r)$. Hence, the convex hull of $B(\bar{V}(q);r)$ fully contains $\bar{V}(q)$, which implies that $\bar{V}(q)$ is contained in $\bar{V}(r)$. □

Proof of Proposition 6. Take two equilibrium paths and any time interval $\tau = [t^1, t^2)$ in which a^1 is played in the first path and a^2 in the second one, and the first jump from either a^1 or a^2 occurs at time instant t^2. Without the loss of generality we may assume that $t^1 = 0$ and $t^2 = 1$. The case when both a^1 and a^2 are played indefinitely long can be handled by taking any arbitrary division of the interval $[t^1, \infty)$ into disjoint subintervals union of which equals $[t^1, \infty)$.

When there are n-players, the original interval $\tau = [0, 1)$ is split into at most $n+1$ pieces determined by $\Delta_i \geq 0$, $i = 1, 2, \ldots, n$. The original interval is decomposed into

$$[0, \Delta_1) \cup [\Delta_1, \Delta_1 + \Delta_2) \cup \cdots \cup \left[\sum_i \Delta_i, 1\right).$$

Let τ_i denote the i'th interval in the decomposition. Note, that there are $n+1$ of this kind of subintervals τ_i. Furthermore, if $\Delta_1, \ldots, \Delta_n = 0$, then the last subinterval equals τ.

The convex combination of the payoffs on the interval τ (ignoring the continuation payoffs) is

$$\lambda(I - e^{-r})u(a^1) + (1-\lambda)(I - e^{-r})u(a^2).$$

The purpose is to show that the interval can be split such that when a^1 is played every second subinterval and a^2 on every other, then the payoffs from the intervals corresponding to a^1 are exactly $\lambda(I - e^{-r})u(a^1)$, and the intervals corresponding to a^2 yield $(1-\lambda)(I - e^{-r})u(a^2)$.

Assume that a^1 is played for any odd index i. Let us set

$$f_i(\Delta_1, \ldots, \Delta_n) = r_i \int_{\substack{t \in \tau_j \\ j \text{ odd}}} e^{-r_i t} u_i(a^1) dt \; i \in N.$$

The resulting function f is continuous, $f(\mathbf{0}) = \mathbf{0}$ for n odd, and $f_i(\mathbf{0}) = (1 - e^{-r_i})u_i(a^1)$ for all i when n is even. Let us now define an auxiliary function $g : S^n \mapsto \mathbb{R}^n$, where S^n is the n-dimensional unit simplex, i.e., $S^n = \{x \in \mathbb{R}^n_+ : \sum_i x_i \leq 1\}$:

$$g_i(\Delta) = \begin{cases} f_i(\Delta) - \lambda(1 - e^{-r_i})u_i(a^1), & \text{if } n \text{ is even} \\ \lambda(1 - e^{-r_i})u_i(a^1) - f_i(\Delta) & \text{if } n \text{ is odd.} \end{cases}$$

When $u_i(a^1) \geq 0$ it holds that $g_i(\mathbf{0}) > 0$ for all $\lambda \in (0, 1)$. Without the loss of generality it can be assumed that $u_i(a^1) \geq 0$ for all $i \in N$, because if not we could just change the sign of $g_i(\Delta)$. Let us next normalize g_i, $i = 1, \ldots, n$, to obtain a function from S^n to itself. First, denote $h(\Delta) = 1 + \sum_i |g_i(\Delta)| > 0$. The normalized function can be written as

$$G_i(\Delta) = \frac{\Delta_i + |g_i(\Delta)|}{h(\Delta)},$$

and evidently G is a function from S^n to itself. Since f_i is continuous, this function is continuous as well. By the Brouwer's fixed-point theorem there is Δ such that $\Delta = G(\Delta)$. This condition on the other hand implies that $|g_i(\Delta)| = 0$ for all $i \in N$. Moreover, $\Delta \neq \mathbf{0}$ by the definition of g. Hence, there is a division of the interval τ, i.e., vector $\Delta \neq \mathbf{0}$ such that $f(\Delta) = \lambda(I - e^{-r})u(a^1)$. Hence, the payoffs corresponding to intervals in which a^1 is played are exactly what we want it to be. On the other hand, $u_i(a^1)$, $i \in N$, are just constants that could have been replaced with $u_i(a^2)$, $i \in N$, without affecting the construction. Hence, the payoffs from the intervals during which a^2 is played are $(1 - e^{-r_i})u_i(a^2) - \lambda(1 - e^{-r_i})u_i(a^2) = (1 - \lambda)(1 - e^{-r_i})u_i(a^2)$, $i \in N$, as supposed to be. Note that the term is the payoff that is obtained if a^2 is played during the whole interval τ, while the second term is the payoff that is not obtained because a^1 is played, the difference is the payoff from those intervals in which a^1 is played.

Because the above construction can be made for all time intervals during which there is no jump in any two equilibrium paths, it is possible to create a convex combination of payoffs of the two paths. This payoff corresponds to a path constructed as above. It remains to be shown that the corresponding path is an equilibrium, i.e., supported by the minimax payoffs. For this purpose it should be shown that there are no profitable deviations from a^1 and a^2 in any interval τ used in the construction of the path.

Let v^1 and v^2 denote the continuation payoffs at time instant $t = 0$ corresponding to the two paths. Observe that the path constructed above yields the players the convex combination of two payoffs that are no less than v_i^-, $i \in N$. Hence, the payoffs $\lambda v_i^1 + (1 - \lambda)v_i^2$ are no less than v_i^-, $i \in N$. Take $t = 0$ to be the beginning of any time interval in which there has been a jump from one action profile to another in either of the two paths. Assume that a^1 and a^2 are played in the time-interval beginning at $t = 0$, when the two paths are followed. If $u_i(a^1), u_i(a^2) \geq v_i^-$ for all $i \in N$, the result is evident. Hence, assume that at least one of $u_i(a^j)$, $i \in N$, $j = 1, 2$, is less than the corresponding v_i^-. Without the loss of generality it can be assumed that $u_i(a^1) < v_i^-$ and $u_i(a^1) < u_i(a^2)$.

In the worst case for player i, the action profile a^1 is played in $[0, t']$, where t' is such that

$$r_i \int_0^{t'} e^{-r\tau} u_i(a^1) d\tau = \lambda(1 - e^{-r_i}) u_i(a^1).$$

However, even in that case player i would rather play a_i^1 than deviate. This is because the largest deviation payoff is obtained by deviating immediately at time instant $t = 0$, which would result to payoff v_i^-, and this payoff is not greater than the payoff from the above path. Hence, there are no profitable one-shot deviations from the constructed path at any time instant. □

References

Abreu, D., Pearce, D., Stacchetti, E., 1986. Optimal cartel equilibria with imperfect monitoring. Journal of Economic Theory 39 (1), 251–269.

Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. Econometrica 58 (5), 1041–1063.

Alós-Ferrer, C., Kern, J., 2015. Repeated games in continuous time as extensive form games. Journal of Mathematical Economics 61, 34–57.

Anderson, R. M., 1984. Quick-response equilibrium. Working Papers in Economic Theory and Econometrics, #IP-323, C.R.M, University of California at Berkeley.

Berg, K., 2018. Set-valued games and mixed-strategy equilibria in discounted supergames. Working paper.

Berg, K., Kitti, M., 2013. Computing equilibria in discounted 2 × 2 supergames. Computational Economics 41 (1), 71–88.

Berg, K., Kitti, M., 2014. Fractal geometry of equilibrium payoffs in discounted supergames. Fractals 22 (4).

Berg, K., Kitti, M., 2016. Equilibrium paths in discounted supergames. Working paper.

Berg, K., Schoenmakers, G., 2017. Construction of subgame-perfect mixed-strategy equilibria in repeated games. Games 8 (4), 47.

Bergin, J., 1992. A model of strategic behavior in repeated games. Journal of Mathematical Economics 21, 113–153.

Bergin, J., 2006. The folk theorem revisited. Economic Theory 27, 321–332.

Bergin, J., MacLeod, W. B., 1992. Continuous time repeated games. International Economic Review 34 (1), 21–37.

Chen, B., Takahashi, S., 2012. A folk theorem for repeated games with unequal discounting. Games and Economic Behavior 76, 571–581.

Cronshaw, M. B., Luenberger, D. G., 1994. Strongly symmetric subgame perfect equilibria in infinitely repeated games with perfect monitoring. Games and Economic Behavior 6 (2), 220–237.

Davidson, R., Harris, R., 1981. Nonconvexities in continuous-time investment. Review of Economic Studies 48, 235–253.

Friedman, A., 1994. Differential games. In: Aumann, R., Hart, S. (Eds.), Handbook of Game Theory, vol 2. Elsevier, Amsterdam, pp. 781–799.

Fudenberg, D., Levine, D., 1986. Limit games and limit equilibria. Journal of Economic Theory 38, 261–279.

Fudenberg, D., Maskin, E., 1986. The folk theorem in repeated games with discounting and incomplete information. Econometrica 54 (3), 533–554.

Kitti, M., 2013. Conditional Markov equilibria in discounted dynamic games. Mathematical Methods of Operations Research 78 (1), 77–100.

Kitti, M., 2016a. Equilibrium payoffs for pure strategies in repeated games. Working paper.

Kitti, M., 2016b. Subgame perfect equilibria in discounted stochastic games. Journal of Mathematical Analysis and Applications 435 (1), 253–266.

Lehrer, E., Pauzner, A., 1999. Repeated games with differential time preferences. Econometrica 67 (2), 393–412.

Murto, P., Välimäki, J., 2013. Delay and information aggregation in stopping games with private information. Journal of Economic Theory 148 (6), 2404–2435.

Rockafellar, R. T., Wets, R. J.-B., 1998. Variational Analysis. Springer, Berlin.

Simon, L., Stinchcombe, M., 1989. Extensive form games in continuous time: pure strategies. Econometrica 57 (5), 1171–1214.

Stinchcombe, M., 1992. Maximal strategy sets for continuous time game theory. Journal of Economic Theory 56 (2), 235–265.

Sugaya, T., 2015. Characterizing the limit set of PPE payoffs with unequal discounting. Theoretical Economics 10, 691–717.

The **Aboa Centre for Economics (ACE)** is a joint initiative of the economics departments of the Turku School of Economics at the University of Turku and the School of Business and Economics at Åbo Akademi University. ACE was founded in 1998. The aim of the Centre is to coordinate research and education related to economics.

Contact information: Aboa Centre for Economics, Department of Economics, Rehtorinpellonkatu 3, FI-20500 Turku, Finland.

www.ace-economics.fi

ISSN 1796-3133