

Subgame-perfection in recursive perfect information games, where each player controls one state

J. Kuipers,^{*}J. Flesch,[†]G. Schoenmakers,^{*} K. Vrieze^{*}

November 28, 2014

Abstract

We consider a class of multi-player games with perfect information and deterministic transitions, where each player controls exactly one non-absorbing state, and where rewards are zero for the non-absorbing states. With respect to the average reward, we provide a combinatorial proof that a subgame-perfect ε -equilibrium exists, for every game in our class and for every $\varepsilon > 0$. We believe that the proof of this result is an important step towards a proof for the more general hypothesis that all perfect information stochastic games, with finite state space and finite action spaces, have a subgame-perfect ε -equilibrium for every $\varepsilon > 0$ with respect to the average reward criterion.

Keywords: perfect information game, recursive game, subgame-perfect equilibrium, average reward.

1 Introduction

We consider a subclass of stochastic games with finite state and action spaces. Shapley [9] introduced the class of zero-sum stochastic games. He proved that such games have a value, under the assumption that there is a positive stopping probability after each move by the players, or equivalently, under the assumption that stage rewards are discounted. Mertens and Neyman [6] demonstrated that every such game also has a value with respect to the average reward. Vieille [13],[14] showed that all two-player stochastic games have an ε -equilibrium for the average reward, for every $\varepsilon > 0$. It follows from the result of Mertens and Neyman in [7] that multi-player stochastic games with deterministic transitions and with perfect information admit an ε -equilibrium for the average reward, for every $\varepsilon > 0$. Thuijsman and

^{*}Dept. of Knowledge Engineering, Maastricht University, P.O. Box 616, 6200MD, Maastricht, The Netherlands.

[†]Dept. of Quantitative Economics, Maastricht University, P.O. Box 616, 6200MD, Maastricht, The Netherlands.

Raghavan [11] showed that they even admit a 0-equilibrium. These ε -equilibria are however not always subgame-perfect. An example by Solan and Vieille [10] demonstrated that a subgame perfect 0-equilibrium need not exist.

The question remains open whether subgame perfect ε -equilibria exist, for every $\varepsilon > 0$. In this paper, we consider a subclass of stochastic games, where each player controls exactly one non-absorbing state, and where non-zero rewards can only be obtained by entering an absorbing state. The technical novelty of this class is that it combines two features that both make it hard to analyse the game: (1) payoffs can be negative in the absorbing states, (2) it may be impossible to move from one non-absorbing state to another non-absorbing state immediately. The technique to deal with these difficulties builds further on those in Flesch et al. [1], Flesch et al. [2], and Kuipers et al. [3], and may be flexible enough to deal with the entire class of perfect information stochastic games.

For recent results on related models, we refer to e.g. Ummels [12], Flesch et al. [2], Purves and Sudderth [8], Flesch et al. [4], and Laraki et al. [5].

2 The class \mathcal{G} and playing a subgame perfect ε -equilibrium

We consider the class \mathcal{G} of games, given by

- (1) a nonempty set of players $N = \{1, \dots, n\}$;
- (2) exactly two states associated with each player $t \in N$: one non-absorbing state identified with t , and one absorbing state denoted by t^* ; the set of absorbing states is denoted by N^* , and the set of all states is denoted by $S = N \cup N^*$;
- (3) for each state $t \in N$, a set of states $A(t) \subseteq N \cup \{t^*\}$ with $t^* \in A(t)$ and $t \notin A(t)$; for each state $t^* \in N^*$, the set $A(t^*)$ is defined as $A(t^*) = \{t^*\}$;
- (4) for each player $t \in N$, an associated (reward) vector $r(t) \in \mathbb{R}^N$.

A game in \mathcal{G} is to be played at stages in \mathbb{N} in the following way. At any stage m one state is called active. If $t \in N$ is active, then player t announces a state in $A(t)$, and the announced state will be active at the next stage. The rewards to the players are zero when this happens. If $t^* \in N^*$ is active, then the unique state $t^* \in A(t^*)$ will be active at the next stage (thus, t^* will be active forever). The stage rewards to the players when this happens are according to $r(t)$, and since $r(t)$ will be the reward at every subsequent stage, $r(t)$ is also the (expected) average reward. The game starts with an initial state $s \in S$.

We assume complete information (i.e. the players know all the data of the game), full monitoring (i.e. the players observe the active state and the action chosen by the active player), and perfect recall (i.e. the players remember the entire sequence of active states and actions).

Playing a game in \mathcal{G} can be interpreted as making an infinite walk in the directed graph $G = (S, E)$, where $E = \{(x, y) \mid x \in S \text{ and } y \in A(x)\}$. In this paper, whenever we refer to an ordered pair (x, y) as an *edge*, it is implicit that $x \in S$ and $y \in A(x)$, i.e. we mean that (x, y) is an element of E .

Let H be a subgraph of G , and denote the edge-set of H by $E(H)$.

Plans. A *plan in H* is an infinite sequence of states $g = (t_m)_{m \in \mathbb{N}}$, such that $(t_m, t_{m+1}) \in E(H)$ for all $t_m \in \mathbb{N}$. A plan in G is simply called a *plan*. A plan is interpreted as a prescription for play for a game with initial active state t_1 . The set of states that become active during play if plan g is executed is denoted by $S(g) \subseteq S$, i.e. $S(g) = \{t \in S \mid \exists m \in \mathbb{N} : t_m = t\}$, and the set of players that become active during play is denoted by $N(g) \subseteq N$, i.e. $N(g) = \{t \in N \mid \exists m \in \mathbb{N} : t_m = t\}$. Notice that, if the initial state of g is an element of N^* , then g is of the form (t^*, t^*, \dots) , with $t^* \in N^*$. Also, if plan g contains a state in N^* , say t^* , and the initial state of g is an element of N , then we must have $t \in N(g)$ and there must be a stage M with $t_M = t$ and with $t_m = t^*$ for all $m > M$. This is interpreted as a prescription for t to announce his absorbing state t^* at stage M . We say that the plan absorbs at t if this is the case. If $S(g) \subseteq N$, then we say that the plan is non-absorbing. We denote by $\phi_t(g)$ the average reward to player t when play is according to g , i.e. $\phi_t(g) = r_t(x)$ if g absorbs at x , and $\phi_t(g) = 0$ if g is non-absorbing. The initial state of plan g is denoted by $\text{FIRST}(g)$.

Paths. A *path (or history) in H* is a finite sequence $p = (t_m)_{m=1}^k$ with $k \geq 1$, such that $(t_m, t_{m+1}) \in E(H)$ for all $m \in \{1, \dots, k-1\}$. A path in G is simply a *path*. The number $k-1$ is called the *length* of the path. The initial state t_1 of path p is denoted by $\text{FIRST}(p)$ and the final state t_k is denoted by $\text{LAST}(p)$. If the length of the path is at least 1, i.e. if the path contains at least one edge, we allow ourselves to say that p is a path *from t_1 to t_k* . We will sometimes want to concatenate a number of paths to make a longer path or a plan, or we may want to concatenate a finite number of paths and a plan to make another plan. We allow concatenation if p^1, p^2, \dots, p^m are paths that satisfy $\text{LAST}(p^k) = \text{FIRST}(p^{k+1})$ for all $k \in \{1, \dots, m-1\}$. The concatenation of these paths is denoted by $\langle p^1, p^2, \dots, p^m \rangle$ and it represents the path that follows the prescription of p^1 from $\text{FIRST}(p^1)$ to $\text{LAST}(p^1) = \text{FIRST}(p^2)$, then follows the prescription of p^2 until $\text{LAST}(p^2) = \text{FIRST}(p^3)$ is reached, and so on, until $\text{LAST}(p^m)$ is reached. Also, if g is a plan with $\text{FIRST}(g) = \text{LAST}(p^m)$, then the plan that first follows the prescription of $\langle p^1, p^2, \dots, p^m \rangle$, and then switches to g is denoted by $\langle p^1, \dots, p^m, g \rangle$. Finally, if we have an infinite number of paths p^1, p^2, \dots with the property $\text{LAST}(p^k) = \text{FIRST}(p^{k+1})$ for all $k \in \mathbb{N}$, then $\langle p^1, p^2, \dots \rangle$ represents the path or plan that subsequently follows the prescription of p^1, p^2 , etc. (The concatenation of an infinite number of paths is still a path if only finitely many of them have positive length.) Let P denote the set of all possible paths, and for $t \in N$, let P^t denote the set of all paths with endpoint t .

Strategies. A *strategy π^t for player t* is a decision rule that, for any path $p \in P^t$, prescribes a probability distribution $\pi^t(p)$ over the elements of $A(t)$. We use the notation Π^t for the set

of strategies for player t . A strategy $\pi^t \in \Pi^t$ is called pure if every prescription $\pi^t(p)$ places probability 1 on one of the elements of $A(t)$. We use the notation Π for the set of joint strategies $\pi = (\pi^t)_{t \in N}$ with $\pi^t \in \Pi^t$ for $t \in N$. A joint strategy $\pi = (\pi^t)_{t \in N}$ is pure if π^t is pure for all $t \in N$.

In this paper, we will define joint strategies by formulating a prescription of play for the entire game, which holds only as long as players execute actions that are assigned a positive probability by the prescription. For the event that a player deviates by choosing an action that is assigned probability 0, a revised prescription will be formulated.

In our description of prescribed play, we distinguish between *default mode* and *threat mode*. Default mode is characterized by deterministic play according to a chosen plan. Threat mode is characterized by play that is deterministic at all stages except one, and where the active player at the exceptional stage is supposed to perform a lottery between exactly two options. The data needed to execute this type of play is given two plans and a player, say by the triple (g, v, x) . Here, g and v are two plans and x is a player, such that $x \in N(g)$, such that $\text{FIRST}(v) \in A(x)$, and such that the state of g following the first occurrence of x differs from $\text{FIRST}(v)$. Threat mode will imply that one of the players in g is active. If this is not player x , or if it is x but not the first occurrence of x on plan g , then the active player is required to follow the prescription of g . If player x is active and if it is his first occurrence on the plan g , then he is required to perform a lottery, in which to place probability ε on the switch to v , and probability $1 - \varepsilon$ on the continuation of g . (The requirement that $\text{FIRST}(v)$ differs from the state following x on g ensures that the lottery places positive probability on two different options.) As long as players follow the plan g , play continues in threat mode according to the triple (g, v, x) . If, on his first occurrence on the plan g , player x switches to v , then play returns to default mode with prescribed plan v . Notice that, in threat mode, player x can deviate from prescribed play, without causing a revision of prescribed play, by performing an alternative lottery for his two options.

Expected rewards. Consider a joint strategy $\pi \in \Pi$ and a path $p \in P$. Suppose that the game has developed along the path p and that state $\text{LAST}(p)$ is now active. Suppose further that all players, starting at $\text{LAST}(p)$, follow the joint strategy π , taking p as the history of the game. Denote the overall probability of absorption at t by $\mathbb{P}^{p, \pi}(t)$. In our model, where non-zero rewards only occur in absorbing states, we see that the expected average reward for player t exists, and that it can be expressed as

$$\psi_t^p(\pi) := \sum_{u \in N} \mathbb{P}^{p, \pi}(u) r_t(u).$$

If π is a pure joint strategy, then following π results in play according to a (deterministic) plan. If we denote this induced plan by $g^{p,\pi}$, then the expected average reward is given by

$$\psi_t^p(\pi) := \sum_{u \in N} \mathbf{P}^{p,\pi}(u) r_t(u) = \phi_t(g^{p,\pi}).$$

If a joint strategy is given in terms of prescribed play, then it is sufficient to know the mode of play (default or threat) at $\text{LAST}(p)$ and the data to describe that mode. If play is in default mode according to plan g , then the expected reward for an arbitrary player $t \in N$ is given by

$$\psi_t^p(\pi) = \phi_t(g).$$

If play is in threat mode according to (g, v, x) , then the expected reward for an arbitrary player $t \in N$ is given by

$$\psi_t^p(\pi) = (1 - \varepsilon)\phi_t(g) + \varepsilon\phi_t(v).$$

Equilibria. Still consider the joint strategy $\pi \in \Pi$ and a game that has developed along the path $p \in P$. The joint strategy $\pi = (\pi^t)_{t \in N} \in \Pi$ is called a (Nash) ε -equilibrium for path p , for some $\varepsilon \geq 0$, if

$$\psi_t^p(\sigma^t, (\pi^u)_{u \in N - \{t\}}) \leq \psi_t^p(\pi) + \varepsilon \quad \forall \sigma^t \in \Pi^t, \forall t \in N,$$

which means that, given history p , no player t can gain more than ε by a unilateral deviation from his proposed strategy π^t to an alternative strategy σ^t . The joint strategy π is called an ε -equilibrium for initial state $s \in N$ if π is an ε -equilibrium for path (s) . The joint strategy π is called a *subgame-perfect ε -equilibrium* if π is an ε -equilibrium for p for every path $p \in P$.

Strategic concepts. We now describe some strategic concepts, necessary to compute and describe a subgame perfect ε -equilibrium for a game in the class \mathcal{G} .

For $\alpha \in \mathbb{R}^N$, a plan g , and a player $x \in N$, we say that x is α -satisfied by g if $\phi_x(g) \geq \alpha_x$. We define

$$\text{SAT}(g, \alpha) = \{x \in N \mid x \text{ is } \alpha\text{-satisfied by } g\}.$$

We say that plan g is α -viable if $N(g) \subseteq \text{SAT}(g, \alpha)$. This means that, if play is according to g , every player t that becomes active during play will receive an average reward of at least α_t . Notice that a plan of the form $g = (t^*, t^*, \dots)$ with $t^* \in N^*$ is trivially α -viable, since $N(g) = \emptyset$. For every state $t \in S$, we denote the set of α -viable plans g with $\text{FIRST}(g) = t$ by $\text{VIABLE}(t, \alpha)$. Notice that, for $t^* \in N^*$, the set $\text{VIABLE}(t^*, \alpha)$ consists of only the plan (t^*, t^*, \dots) .

For $t \in N$, and a plan $g \in \text{VIABLE}(u, \alpha)$ with $u \in A(t)$, say that g is (t, u, α) -admissible if at least one of the following holds:

AD-i: $t \notin N(g)$ or g is non-absorbing;

AD-ii: $\alpha_t > 0$;

AD-iii: $t \in N(g)$ and there exists a pair (x, v) , such that $x \neq t$ is a player who resides on g in the part from u to the first occurrence of t , such that v is an α -viable plan with $\text{FIRST}(v) \in A(x)$, such that $\text{FIRST}(v)$ is not the state following the first occurrence of x on plan g , and such that $x, t \notin \text{SAT}(v, \alpha)$.

We denote the set of plans that are (t, u, α) -admissible by $\text{ADMISS}(t, u, \alpha)$.

The motivation behind the definition of (t, u, α) -admissibility is that player t will prefer (or cannot prevent) to follow through with a prescribed plan $g \in \text{ADMISS}(t, u, \alpha)$, as compared to deviating to u every time he is active, and thereby creating a non-absorbing plan. Condition *AD-i* describes the situation, where the announcement of u by player t is not possible (the case $t \notin N(g)$) or would yield the same average reward as g (the case that g is non-absorbing). Condition *AD-ii* describes the situation, where this deviation, if it is possible, would yield a strictly lower average reward for t , as the non-absorbing plan yields zero average reward, while g yields an average reward of at least $\alpha_t > 0$. Condition *AD-iii* describes the situation, where a player x with $x \neq t$ and participating in the non-absorbing plan created by t 's deviation to u , has the threat of switching to an α -viable plan v with $t, x \notin \text{SAT}(v, \alpha)$. Now, if prescribed play is not exactly g , but if player x is required to place a very small probability on a switch to v , every time he is active, then play will still be according to g with very high probability if players follow the prescription. Player t will not be able to profit from continuous deviation to u , as this would eventually result in a switch to v with probability 1 and $t \notin \text{SAT}(v, \alpha)$. Player x will not be tempted to increase the probability of a switch to v , since $x \notin \text{SAT}(v, \alpha)$.



Figure 1

Example 2.1. Figure 1 represents two games, each with two players, 1 and 2. Only the two nodes that correspond to the non-absorbing states of these players are depicted, and the vectors below these states are the rewards that will be obtained if the controlling player chooses his absorbing state. For both games, the arc from 2 to 1 indicates that $1 \in A(2)$, hence $A(2) = \{1, 2^*\}$, and the other arc indicates $2 \in A(1)$, hence $A(1) = \{2, 1^*\}$.

For game I, let us take $\alpha_t = r_t(t)$ for $t = 1, 2$, i.e. $\alpha_1 = 1$ and $\alpha_2 = -1$. Then, every plan that absorbs at 2 is α -viable. The plan $(1, 1^*, 1^*, \dots)$ is also α -viable, but any plan in which player

2 is active and that absorbs at 1 is not α -viable, since such a plan does not α -satisfy player 2. The two non-absorbing plans $(1, 2, 1, 2, \dots)$ and $(2, 1, 2, 1, \dots)$ are not α -viable, since player 1 is active in both plans and is not α -satisfied by them.

The plans in $\text{VIABLE}(2, \alpha)$ that do not immediately absorb at 2 are not $(1, 2, \alpha)$ -admissible, since they violate all three conditions $AD-i$, $AD-ii$, and $AD-iii$. The plan $(2, 2^*, 2^*, \dots)$ is $(1, 2, \alpha)$ -admissible, since it satisfies $AD-i$. All plans in $\text{VIABLE}(1, \alpha)$ are also $(2, 1, \alpha)$ -admissible, since these plans satisfy $AD-iii$, where player 1 is in the role of 'threat-player' x and plan $(1^*, 1^*, \dots)$ is in the role of 'threat-plan' v .

For game II, let us also take $\alpha_t = r_t(t)$ for $t = 1, 2$, i.e. $\alpha_1 = \alpha_2 = -2$. Here, it is easily verified that every plan is α -viable. Any plan that starts at 1 and absorbs at 2 is not $(2, 1, \alpha)$ -admissible however, since conditions $AD-i$, $AD-ii$, and $AD-iii$ are all three violated for such plans. The plans $(1, 1^*, 1^*, \dots)$ and $(1, 2, 1, 2, \dots)$ are $(2, 1, \alpha)$ -admissible, as they both satisfy $AD-i$. By symmetry of the example, any plan that starts at 2 and absorbs at 1 is not $(1, 2, \alpha)$ -admissible, but the plans $(2, 2^*, 2^*, \dots)$ and $(2, 1, 2, 1, \dots)$ are. \diamond

Update procedure. We now propose a method for updating one coordinate of a vector $\alpha \in \mathbb{R}^N$. For $t \in N$, $u \in A(t)$, and $\alpha \in \mathbb{R}^N$, we define

$$\beta(t, u, \alpha) = \min\{\phi_t(v) \mid v \in \text{ADMISS}(t, u, \alpha)\}.$$

We define further,

$$\delta(t, \alpha) = \max\{\beta(t, u, \alpha) \mid u \in A(t)\},$$

and

$$B(t, \alpha) = \{u \in A(t) \mid \beta(t, u, \alpha) = \delta(t, \alpha)\}.$$

We use the convention $\min \emptyset = \infty$, so that $\beta(t, u, \alpha)$ and $\delta(t, \alpha)$ are well defined. Moreover, by the definition of $\delta(t, \alpha)$, the set $B(t, \alpha)$ is always nonempty.

Example 2.2. Let us apply the updating procedure to the two games depicted in figure 1, with the initial α -values from example 2.1.

For game I, we chose $\alpha_1 = 1$ and $\alpha_2 = -1$. Let us first update the state controlled by player 1. We trivially have $\beta(1, 1^*, \alpha) = r_1(1) = 1$ and we have $\beta(1, 2, \alpha) = \phi_1(2, 2^*, 2^*, \dots) = r_1(2) = 2$, since $(2, 2^*, 2^*, \dots)$ is the only $(1, 2, \alpha)$ -admissible plan. Thus, $\delta(1, \alpha) = \max(1, 2) = 2$, and the updated vector is given by $\alpha_1^* = 2$ and $\alpha_2^* = -1$. Let us now update the updated vector α^* . This time, we update the state controlled by player 2. We trivially have $\beta(2, 2^*, \alpha) = r_2(2) = -1$ and one can verify that $\beta(2, 1, \alpha) = \phi_1(1, 2, 2^*, 2^*, \dots) = r_2(2) = -1$. Thus, $\delta(2, \alpha) = \max(-1, -1) = -1$, and we observe that this update had no effect. One can verify that an update of the already updated state controlled by 1 will have no further effect either.

For game II, we chose $\alpha_1 = \alpha_2 = -2$. Let us denote $\alpha^0 = \alpha$, and for $i > 0$, let us use the notation α^i for the vector that results after an update of α^{i-1} . We choose to update first the state

controlled by player 1. We trivially have $\beta(1, 1^*, \alpha) = r_1(1) = -2$, and the findings in example 2.1 show that $\beta(1, 2, \alpha) = \min(\phi_1(2, 1, 2, 1, \dots), \phi_1(2, 2^*, 2^*, \dots)) = -1$. Thus, $\delta(1, \alpha) = -1$ and $\alpha^1 = (-1, -2)$. We now set $\alpha = \alpha^1$ and we continue with an update of the state controlled by player 2. We have $\beta(2, 2^*, \alpha) = r_2(2) = -2$. Notice that $(1, 2, 1, 2, \dots)$ is the only $(2, 1, \alpha)$ -admissible plan. (The plan $(1, 1^*, 1^*, \dots)$ is not $(2, 1, \alpha)$ -admissible 'anymore', due to the previous update.) Thus, we have $\beta(2, 1, \alpha) = \phi_1(1, 2, 1, 2, \dots) = 0$. It follows that $\alpha^2 = (-1, 0)$. We set $\alpha = \alpha^2$. Notice that only the two non-absorbing plans $(1, 2, 1, 2, \dots)$ and $(2, 1, 2, 1, \dots)$ are now α -viable. They are also respectively $(2, 1, \alpha)$ -admissible and $(1, 2, \alpha)$ -admissible. Then $\beta(2, 1, \alpha) = \phi_1(1, 2, 1, 2, \dots) = 0$, which means that we obtain $\alpha^3 = (0, 0)$ by updating the state controlled by 1 for the second time. After this, no update will have any further effect. \diamond

Main result. In Sections 3 and 4, we analyze the update procedure. The analysis demonstrates that, starting with an appropriate initial vector $\alpha^0 \in \mathbb{R}^N$, repeated application of the update procedure will produce, in a finite number of iterations, a vector $\alpha^* \in \mathbb{R}^N$, such that $\delta(t, \alpha^*) = \alpha_t^*$ for all $t \in N$. The existence of such a 'fixed point' α^* is proven in theorem 4.16 and it allows for the construction of a subgame-perfect ε -equilibrium for all $\varepsilon > 0$. The following theorem is our main result.

Theorem 2.1. *There exists a subgame perfect ε -equilibrium for every game in class \mathcal{G} and every $\varepsilon > 0$.*

Proof. I: Description of a joint strategy. Take a game in \mathcal{G} . As we remarked, a vector $\alpha^* \in \mathbb{R}^N$ exists such that $\delta(t, \alpha^*) = \alpha_t^*$ for all $t \in N$. The fact that $\delta(t, \alpha^*)$ is finite for all $t \in N$ implies that $\beta(t, u, \alpha^*)$ is finite for all $t \in N$ and all $u \in A(t)$. It follows that $\text{VIABLE}(t, \alpha^*)$ is non-empty for all $t \in N$ and that $\text{ADMISS}(t, u, \alpha^*)$ is non-empty for all $t \in N$ and $u \in A(t)$.

We can thus choose a plan $g^t \in \text{VIABLE}(t, \alpha^*)$ for all $t \in N$. We can also choose, for all $t \in N$ and all $u \in A(t)$, a plan $g^{tu} \in \text{ADMISS}(t, u, \alpha^*)$ such that $\phi_t(g^{tu}) = \beta(t, u, \alpha^*)$. In case the choice g^{tu} violates both *AD-i* and *AD-ii*, and therefore satisfies *AD-iii*, we can make the additional choice of a player $x^{tu} \in N \setminus \{t\}$ and an α^* -viable plan v^{tu} , such that $\text{FIRST}(v^{tu}) \in A(x^{tu})$, such that $\text{FIRST}(v)$ is not the state following the first occurrence of x^{tu} on plan g^{tu} , such that x^{tu} resides on the non-absorbing plan that would result if player t were to deviate from g^{tu} by announcing u every time he is active, and such that $x^{tu}, t \notin \text{SAT}(v^{tu}, \alpha^*)$.

Now, if $s \in N$ is the initial state of the game, then we start the game in default mode with plan g^s as initial prescription. Whenever a deviation from prescribed play is detected, say that player t deviates to u , then we check if plan g^{tu} , associated with the pair (t, u) , satisfies *AD-i* or *AD-ii*. If so, then revised play will be in default mode according to the (newly) prescribed plan g^{tu} . If not, then *AD-iii* is satisfied, and play will resume in threat mode according to the triple (g^{tu}, v^{tu}, x^{tu}) . To be clear, if play was already in threat mode according to the same triple (g^{tu}, v^{tu}, x^{tu}) , then plan g^{tu} is restarted at its initial state. Let us denote the joint strategy defined in this way by π .

II: Verification. Let us verify that π is a subgame perfect $2\varepsilon M$ -equilibrium, where $M = \max\{|r_i(t)| \mid i, t \in N\}$. Assume that play has developed along path p . Let us further assume that exactly one player deviates from π , after the game has reached $\text{LAST}(p)$. We will demonstrate that the deviating player can profit at most $2\varepsilon M$ in expectation from his deviation(s).

IIa: First assume that a player, say y , deviates a finite number of times after the game has developed along p . Say that at his last deviation he announces $u \in A(y)$, and that this last deviation causes a revision of prescribed play.

If the plan g^{yu} satisfies *AD-i* or *AD-ii*, then after y 's last deviation, play will resume in default mode with prescription g^{yu} . The plan g^{yu} was selected such that $\phi_y(g^{yu}) = \beta(y, u, \alpha^*)$. By definition of $\delta(y, \alpha^*)$, we have $\beta(y, u, \alpha^*) \leq \delta(y, \alpha^*)$, and by the properties of α_y^* , we have $\delta(y, \alpha^*) = \alpha_y^*$. It follows that $\phi_y(g^{yu}) \leq \alpha_y^*$. If the plan g^{yu} violates both *AD-i* and *AD-ii*, then after y 's last deviation, play will resume in threat mode according to the triple (g^{yu}, v^{yu}, x^{yu}) . Since no further deviations will take place, plan g^{yu} will be executed with probability $1 - \varepsilon$ and with an expected average reward of at most α_y^* for player y ; plan v^{yu} will be executed with probability ε and with an expected average reward strictly less than α_y^* for player y , since $y \notin \text{SAT}(v^{yu}, \alpha^*)$. Thus, the expected average reward for y , after deviations, is at most $(1 - \varepsilon)\alpha_y^* + \varepsilon\alpha_y^* = \alpha_y^*$.

Let us now demonstrate that player y has an expected average reward of at least $\alpha_y^* - 2\varepsilon M$ if he follows π . If y becomes first active during default mode, it means that prescribed play is deterministic according to an α^* -viable plan, say g . Then y 's average reward will be $\phi_y(g)$ and since $y \in N(g)$, we have indeed $\psi_y^p(\pi) = \phi_y(g) \geq \alpha_y^* > \alpha_y^* - 2\varepsilon M$. If player y becomes active during threat mode, then play is according one of the plans g^{tu} . If the first occurrence of player x^{tu} comes before y on plan g^{tu} , then there is no chance of a switch to v^{tu} anymore, and the average reward for player y will be $\phi_y(g^{tu}) \geq \alpha_y^* > \alpha_y^* - 2\varepsilon M$. If the first occurrence of x^{tu} comes after y or if $x^{tu} = y$, then there is probability $1 - \varepsilon$ that plan g^{tu} will be executed, with an average reward of at least α_y^* for player y . There is also a probability of ε that v^{tu} will be executed, with an average reward of at least $-M$ for player y . Thus indeed, the expected average reward for player y is $\psi_y^p(\pi) \geq (1 - \varepsilon)\alpha_y^* - \varepsilon M \geq \alpha_y^* - 2\varepsilon M$.

IIb: Now assume that a player, say y , deviates a finite number of times after the game has developed along p , such that his last deviation does *not* cause a revision of prescribed play. Then, at the last deviation of y , the game was in threat mode and player y was the one to perform the lottery. Say that the game was in threat mode according to the triple (g, v, y) . Notice that the player to cause a revision of play is never the one assigned the task of performing a lottery in revised play. Therefore, player y could not have deviated in the period between the stage where the game had developed along p and the stage at which he was supposed to perform the lottery. So, the last deviation of y was in fact his only deviation in the relevant time period, and threat mode according to (g, v, y) was the mode at $\text{FIRST}(p)$. We conclude

that $\psi_y^p(\pi) = (1 - \varepsilon)\phi_y(g) + \varepsilon\phi_y(v)$ is the expected average reward for y by following π .

Now say that in the alternative lottery performed by player y , he places probability μ on the continuation of g , probability $\lambda \leq 1 - \mu$ on the switch to v , and summed probability $\eta = 1 - \mu - \lambda$ on his other available actions. Then there is probability μ that g is executed, with expected average reward $\phi_y(g)$, probability λ that v is executed, with expected average reward strictly less than α_y^* , since $y \notin \text{SAT}(v, \alpha^*)$, and probability η that his deviation will be detected, with expected average reward at most α_y^* . His expected average reward is then at most $\mu\phi_y(g) + \lambda\alpha_y^* + \eta\alpha_y^* \leq \phi_y(g)$, where the inequality is due to the fact that $\phi_y(g) \geq \alpha_y^*$. Since the expected average reward for following π is $\psi_y^p(\pi) = (1 - \varepsilon)\phi_y(g) + \varepsilon\phi_y(v) \geq \phi_y(g) - 2\varepsilon M$, we see that player y profits no more than $2\varepsilon M$ from his deviation.

IIc: It remains to investigate the possibility that a player deviates infinitely many times. As this implies infinite play along non-absorbing states, the average reward to all players will be zero if this happens. If the deviating player, say t , satisfies $\alpha_t^* \geq 0$, then he will profit at most $2\varepsilon M$, as he will receive at least $\alpha_t^* - 2\varepsilon M$ by sticking to the plan. So we can assume that $\alpha_t^* < 0$. Notice that, if t deviates infinitely many times, he also causes infinitely many times a revision of prescribed play. Each time at such a revision, it will be checked if the plan g^{tu} satisfies *AD-i* or *AD-ii*, where u is the state to which t deviates. By the choice of g^{tu} as a minimizer for $\beta(t, u, \alpha^*)$, we have $\phi_t(g^{tu}) = \beta(t, u, \alpha^*) \leq \delta(t, \alpha^*) = \alpha_t^* < 0$. This implies that g^{tu} is an absorbing plan. Also, since t will deviate again, we must have $t \in \mathbb{N}(g^{tu})$. This means that *AD-i* is violated, and clearly *AD-ii* is violated too, by our assumption $\alpha_t^* < 0$. Thus, only *AD-iii* is satisfied, and play will resume in threat mode according to the triple (g^{tu}, v^{tu}, x^{tu}) .

We conclude that, in the event of infinitely many deviations by a player t with $\alpha_t^* < 0$, after a while, play will be constantly in threat mode. From then on, between every two visits of t during play, a probability ε is placed on a switch to one of the plans v^{tu} . After such a switch, player t will not become active again, since $t \notin \text{SAT}(v^{tu}, \alpha^*)$, and therefore $t \notin \mathbb{N}(v^{tu})$. Thus, once play gets in constant threat mode, there is a probability of at most $(1 - \varepsilon)^n$ that there will be at least n more visits to player t . It follows that the event of infinitely many deviations by a player t with $\alpha_t^* < 0$ has probability 0. \square

Example 2.3. For game I, a fixed point $\alpha_1^* = 2$ and $\alpha_2^* = -1$ was calculated in example 2.2. All α^* -viable plans absorb at 2, so we define $g^1 = g^{21} = (1, 2, 2^*, 2^*, \dots)$ and $g^2 = g^{12} = (2, 2^*, 2^*, \dots)$. Plan g^{12} is $(1, 2, \alpha^*)$ -admissible, since it satisfies *AD-i*. Plan g^{21} is $(2, 1, \alpha^*)$ -admissible, since it satisfies *AD-iii*, but not *AD-i* and *AD-ii*. Then, in order to complete our choices, we choose player 1 in the role of 'threat-player' x^{12} and plan $(1^*, 1^*, \dots)$ in the role of 'threat-plan' $v^{12} = (1^*, 1^*, \dots)$.

With these choices, the proof of theorem 2.1 now prescribes absorption at 2 whenever player 2 is active (execute either g^2 or g^{12}). If the game starts at 1, then player 1 is initially supposed to announce 2 with probability 1 (execute g^1). Every next time he becomes active, player 1 is supposed to announce 2 (execute g^{12}) with probability $1 - \varepsilon$ and to absorb at 1 (execute plan

v^{12}) with probability ε . (Player 2 apparently refused to absorb at 2, and play is now in threat mode.)

For game II, the fixed point is $\alpha_1^* = 0$ and $\alpha_2^* = 0$. Here, only the two non-absorbing plans $(1, 2, 1, 2, \dots)$ and $(2, 1, 2, 1, \dots)$ are α^* -viable. We therefore choose $g^1 = (1, 2, 1, 2, \dots)$ and $g^2 = (2, 1, 2, 1, \dots)$. We do not make any further choices. (Since it is not a deviation to announce the other player, it is unnecessary to choose the plans g^{21} and g^{12} .)

With these choices, the proof of theorem 2.1 now prescribes the players to always announce each other, indefinitely. Deviation from the plan means absorption, which is indeed not profitable for the deviating player. \diamond

Of the restrictions imposed on our class \mathcal{G} of games, the requirement that each player control just one non-absorbing state seems especially severe. The following example is to demonstrate that at least one specific attempt to deal with multiple states per player does not work.




Figure 2

Example 2.4. The update procedure could be applied to the 1-player game in figure 2, as if the two states were controlled by two different players, say 1 and $1'$. If we initiate the update procedure with $\alpha_t = -1$ for $t = 1, 1'$, then we find that α is in fact a fixed point. The construction of theorem 2.1 then allows, for any $\varepsilon > 0$, for the strategy profile in which player 1 absorbs immediately with probability 1, in any subgame. So clearly, the construction of theorem 2.1 fails here. \diamond

3 Semi-stable vectors and their properties

The purpose of this section is to provide a sufficient condition for $\alpha \in \mathbb{R}^N$ that guarantees the existence of a (t, u, α) -admissible plan for every $t \in N$ and every $u \in A(t)$.

For $\alpha \in \mathbb{R}^N$, $t \in N$ and $u \in S$, let us say that t is α -safe at u if $u \in A(t)$ and if $t \in \text{SAT}(g, \alpha)$ for all $g \in \text{VIABLE}(u, \alpha)$. For $t^* \in N^*$, it will be convenient to say that t^* is α -safe at t^* . We define, for all $t \in S$,

$$\text{SAFESTEP}(t, \alpha) = \{u \in A(t) \mid t \text{ is } \alpha\text{-safe at } u\}.$$

For $\alpha \in \mathbb{R}^N$ and $X \subseteq N$, we define

$$\text{POS}(X, \alpha) = \{x \in X \mid \alpha_x > 0\},$$

$$\text{ESC}(X, \alpha) = \{x \in X \mid \exists y \in S \setminus X : y \in \text{SAFESTEP}(x, \alpha)\}.$$

We also define

$$\mathcal{C} = \{X \subseteq N \mid A(x) \cap X \neq \emptyset \text{ for all } x \in X\},$$

and for $\alpha \in \mathbb{R}^N$, we define

$$\mathcal{P}(\alpha) = \{X \subseteq N \mid \text{POS}(X, \alpha) \neq \emptyset\},$$

$$\mathcal{E}(\alpha) = \{X \subseteq N \mid \text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha) = \emptyset\}, \text{ and}$$

$$\mathcal{X}(\alpha) = \mathcal{P}(\alpha) \cap \mathcal{E}(\alpha) \cap \mathcal{C}.$$

We say that an edge $e = (x, y)$ is an α -exit from X if $x \in X$ and $y \in S \setminus X$, and if, for all $g \in \text{VIABLE}(y, \alpha)$,

$$\text{ESC}(X, \alpha) \subseteq \text{SAT}(g, \alpha) \implies x \in \text{SAT}(g, \alpha).$$

We say that e is a trivial α -exit from X if $x \in \text{ESC}(X, \alpha)$ and a non-trivial one if $x \in X \setminus \text{ESC}(X, \alpha)$.

We now say that $\alpha \in \mathbb{R}^N$ is *semi-stable* if $\text{SAFESTEP}(x, \alpha) \neq \emptyset$ for all $x \in N$, and if there exists a non-trivial α -exit from X for every $X \in \mathcal{X}(\alpha)$. The set of semi-stable vectors in \mathbb{R}^N is denoted by Ω .

Lemma 3.1. *We have $\rho \in \Omega$, where ρ is the vector defined by $\rho_t = r_t(t)$ for all $t \in N$.*

Proof. For the vector ρ , we have $t^* \in \text{SAFESTEP}(t, \rho)$ for all $t \in N$. Hence, $\text{SAFESTEP}(t, \rho) \neq \emptyset$ for all $t \in N$.

For any $X \subseteq N$, we have $\text{ESC}(X, \rho) = X$, since $t^* \in \text{SAFESTEP}(t, \rho)$ for all $t \in X$. Now, let $X \in \mathcal{P}(\rho)$, i.e. let $X \subseteq N$ be such that $\text{POS}(X, \rho) \neq \emptyset$. Then $\text{ESC}(X, \rho) \cap \text{POS}(X, \rho) = \text{POS}(X, \rho) \neq \emptyset$, hence $X \notin \mathcal{E}(\rho)$. It follows that $\mathcal{X}(\rho) = \mathcal{P}(\rho) \cap \mathcal{E}(\rho) \cap \mathcal{C} = \emptyset$. Then trivially, a non-trivial ρ -exit from X exists for all $X \in \mathcal{X}(\rho)$. Thus, $\rho \in \Omega$. \square

Lemma 3.2. *Let $\alpha \in \Omega$. Then $\text{ESC}(X, \alpha) \neq \emptyset$ for all $X \in \mathcal{P}(\alpha)$.*

Proof. Let $X \in \mathcal{P}(\alpha)$ and suppose that $\text{ESC}(X, \alpha) = \emptyset$. Then $\text{SAFESTEP}(x, \alpha) \subseteq X$ for all $x \in X$, and thus, $\text{SAFESTEP}(x, \alpha) \subseteq A(x) \cap X$ for all $x \in X$. We further have $\text{SAFESTEP}(x, \alpha) \neq \emptyset$ for all $x \in X$, by the fact that $\alpha \in \Omega$. It follows that $A(x) \cap X \supseteq \text{SAFESTEP}(x, \alpha) \neq \emptyset$ for all $x \in X$, which proves that $X \in \mathcal{C}$. We also trivially have $X \in \mathcal{E}(\alpha)$, since we suppose $\text{ESC}(X, \alpha) = \emptyset$.

Then $X \in \mathcal{P}(\alpha) \cap \mathcal{E}(\alpha) \cap \mathcal{C} = \mathcal{X}(\alpha)$. Now, since $\alpha \in \Omega$ and $X \in \mathcal{X}(\alpha)$, there exists a non-trivial α -exit from X , say (x, y) . Since (x, y) is an α -exit from X , we have for all $v \in \text{VIABLE}(y, \alpha)$,

$$\text{ESC}(X, \alpha) \subseteq \text{SAT}(v, \alpha) \implies x \in \text{SAT}(v, \alpha).$$

The leftside of the implication is always true, since supposedly $\text{ESC}(X, \alpha) = \emptyset$. It follows that $x \in \text{SAT}(v, \alpha)$ for all $v \in \text{VIABLE}(y, \alpha)$, hence $y \in \text{SAFESTEP}(x, \alpha)$. But this implies $x \in \text{ESC}(X, \alpha)$, which contradicts that $\text{ESC}(X, \alpha) = \emptyset$. \square

For any subgraph H of G and a subset X of the vertex set $V(H)$ of H , say that X is an *ergodic set* of H if

- i) for all $x, y \in X$, there exists a path p in H from x to y with $N(p) \subseteq X$, and
- ii) for all $x \in X$ and $y \in V(H) \setminus X$, there is no path in H from x to y .

The following lemma is an easy result in graph theory. It is stated without proof.

Lemma 3.3. *Let $H = (V(H), E(H))$ be a (directed) graph, such that for every vertex $x \in V(H)$, there exists $y \in V(H)$ with $(x, y) \in E(H)$. Then, for every $x \in V(H)$, there is a path from x to an element of an ergodic set of H .*

For $\alpha \in \Omega$, define the graph $G(\alpha)$ as the graph with vertex set S and edge set $\{(x, y) \in E \mid y \in \text{SAFESTEP}(x, \alpha)\}$. Notice that, for all $\alpha \in \Omega$ and for all $t^* \in N^*$, the singleton $\{t^*\}$ is an ergodic set of the graph $G(\alpha)$, since (t^*, t^*) is a path in $G(\alpha)$ from t^* to t^* and since there is no edge leaving the set $\{t^*\}$. The definition of ergodic set implies that different ergodic sets of a graph are disjoint. Therefore, any ergodic set of $G(\alpha)$ is either a singleton from the set N^* or a subset of N . The following corollary follows directly from lemma 3.2.

Corollary 3.4. *Let $\alpha \in \Omega$. If $X \subseteq N$ is an ergodic set of $G(\alpha)$, then $\text{POS}(X, \alpha) = \emptyset$.*

An immediate insight from corollary 3.4 is that, for all $\alpha \in \Omega$, a plan v is always an α -viable plan if $N(v)$ is a subset of an ergodic set of $G(\alpha)$.

Lemma 3.5. *Let $\alpha \in \Omega$, let p be a path in $G(\alpha)$, and let g be an α -viable plan such that $\text{FIRST}(g) = \text{LAST}(p)$. Then the plan $\langle p, g \rangle$ is α -viable.*

Proof. Write $p = (z_i)_{i=1}^k$ with $k \geq 1$. Define, for all $i \in \{1, \dots, k\}$, the plan $g_i = \langle (z_i, \dots, z_k), v \rangle$. We prove by induction on i that all plans g_i with $i \in \{1, \dots, k\}$ are α -viable. Trivially, the plan $g_k = g$ is α -viable. Now assume that g_{i+1} is α -viable with $i < k$. Then $N(g_i) \setminus \{z_i\} \subseteq N(g_{i+1}) \subseteq \text{SAT}(g_{i+1}, \alpha) = \text{SAT}(g_i, \alpha)$. Thus, to prove that $N(g_i) \subseteq \text{SAT}(g_i, \alpha)$, i.e. to prove that g_i is α -viable, it suffices to show that $z_i \in \text{SAT}(g_i, \alpha)$.

We have $z_{i+1} \in \text{SAFESTEP}(z_i, \alpha)$, since p is a path in $G(\alpha)$. It follows that $z_i \in \text{SAT}(g_{i+1}, \alpha)$, since g_{i+1} is an α -viable plan with $\text{FIRST}(g_{i+1}) = z_{i+1}$. Because $\text{SAT}(g_i, \alpha) = \text{SAT}(g_{i+1}, \alpha)$, we indeed obtain $z_i \in \text{SAT}(g_i, \alpha)$, \square

The existence of α -viable plans for $\alpha \in \Omega$ is now an easy consequence of previous results.

Lemma 3.6. *Let $\alpha \in \Omega$. Then, for all $t \in N$, a plan g in $G(\alpha)$ exists with $g \in \text{VIABLE}(t, \alpha)$.*

Proof. Let $t \in N$. We have $\text{SAFESTEP}(x, \alpha) \neq \emptyset$ for all $x \in N$, since $\alpha \in \Omega$. Thus, the graph $G(\alpha)$ satisfies the conditions of lemma 3.3. Therefore, there is a path p in $G(\alpha)$ from t to an element x of an ergodic set X of $G(\alpha)$. By the properties of an ergodic set, a path q in $G(\alpha)$ exists from x to x , and with $N(q) \subseteq X$. First we prove that the plan $\langle q, q, \dots \rangle$ is α -viable. Indeed, this is true by definition if the plan $\langle q, q, \dots \rangle$ is absorbing, i.e. if $x \in N^*$ and $q = (x, x)$. Otherwise, we have $X \subseteq N$, and $\text{POS}(X, \alpha) = \emptyset$ follows by corollary 3.4. The plan $\langle q, q, \dots \rangle$ is then α -viable, due to the fact that it is non-absorbing, hence it gives average reward 0 to all players. Finally, it follows that $g := \langle p, q, q, \dots \rangle \in \text{VIABLE}(t, \alpha)$, by lemma 3.5. \square

The main result of this section concerns the existence of admissible plans.

Lemma 3.7. *Let $\alpha \in \Omega$. Then, for all $t \in N$ and for all $u \in A(t)$, $\text{ADMISS}(t, u, \alpha) \neq \emptyset$.*

Proof. Let $t \in N$ and $u \in A(t)$. By lemma 3.6, $\text{VIABLE}(u, \alpha) \neq \emptyset$. If $v \in \text{VIABLE}(u, \alpha)$ exists with $t \notin N(v)$, then $v \in \text{ADMISS}(t, u, \alpha)$ since v satisfies *AD-i*, and we are done. Assume further that $t \in N(v)$ for all $v \in \text{VIABLE}(u, \alpha)$. Notice that this implies $t \in \text{SAT}(v, \alpha)$ for all $v \in \text{VIABLE}(u, \alpha)$. Thus, $u \in \text{SAFESTEP}(t, \alpha)$, i.e. (t, u) is an edge of $G(\alpha)$.

Define

$$Y = \{y \in S \mid \text{a path } p \text{ in } G(\alpha) \text{ from } u \text{ to } y \text{ exists with } t \notin N(p)\},$$

and define $X = Y \cup \{u, t\}$. We claim that $\text{ESC}(X, \alpha) \subseteq \{t\}$. To prove our claim, we let $x \in X \setminus \{t\}$, and we will show that $y \in X$ for all $y \in \text{SAFESTEP}(x, \alpha)$. So, let $y \in \text{SAFESTEP}(x, \alpha)$. If $y = t$, then trivially $y \in X$. If $y \neq t$, then there is a path in $G(\alpha)$ from u to y not containing t . Indeed, if $x = u$, then $(x, y) = (u, y)$ is such a path, and if $x \neq u$, then there is a path p in $G(\alpha)$ from u to x not containing t by the fact that $x \in X \setminus \{t, u\} \subseteq Y$, and $\langle p, (x, y) \rangle$ is then a path in $G(\alpha)$ from u to y not containing t . Thus, $y \in Y \subseteq X$ as claimed. We now distinguish between the cases $\text{POS}(X, \alpha) = \emptyset$ and $\text{POS}(X, \alpha) \neq \emptyset$.

First assume that $\text{POS}(X, \alpha) = \emptyset$. Notice that, for all $x \in X$, an element $y \in X$ exists, such that (x, y) is an edge of $G(\alpha)$. In particular also, (t, u) is an edge of $G(\alpha)$. Then it is possible to construct a non-absorbing plan g with $\text{FIRST}(g) = u$, with $N(g) \subseteq X$, and such that every edge of g is in the edge-set of $G(\alpha)$. Then $g \in \text{VIABLE}(u, \alpha)$ by the assumption that $\text{POS}(X, \alpha) = \emptyset$, and by the fact that a non-absorbing plan gives reward 0 to all players. Since the non-absorbing plan g satisfies condition *AD-i*, it also follows that $g \in \text{ADMISS}(t, u, \alpha)$.

Now assume that $\text{POS}(X, \alpha) \neq \emptyset$, i.e. assume that $X \in \mathcal{P}(\alpha)$. We then have $\text{ESC}(X, \alpha) \neq \emptyset$ by lemma 3.2, so we must have $\text{ESC}(X, \alpha) = \{t\}$. If $t \in \text{POS}(X, \alpha)$, then $\text{ADMISS}(t, u, \alpha) = \text{VIABLE}(u, \alpha) \neq \emptyset$, where the equality is by the fact that *AD-ii* is satisfied by all plans in $\text{VIABLE}(u, \alpha)$ and the inequality is by lemma 3.6. So we can further assume that $t \notin \text{POS}(X, \alpha)$. Under this assumption, we have $\text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha) = \emptyset$, i.e. $X \in \mathcal{E}(\alpha)$. We also have that $A(x) \cap X \neq \emptyset$ for all $x \in X$, which follows from our earlier observation that, for all $x \in X$, an element $y \in X$ exists, such that (x, y) is an edge of $G(\alpha)$. Thus, we have $X \in \mathcal{C}$.

We see that $X \in \mathcal{P}(\alpha) \cap \mathcal{E}(\alpha) \cap \mathcal{C} = \mathcal{X}(\alpha)$. Then, by the fact that $\alpha \in \Omega$, a non-trivial α -exit from X exists, say (x, y) . By definition of an α -exit, we have, for all $g \in \text{VIABLE}(y, \alpha)$,

$$\text{ESC}(X, \alpha) \subseteq \text{SAT}(g, \alpha) \implies x \in \text{SAT}(g, \alpha).$$

Since $\text{ESC}(X, \alpha) = \{t\}$, this translates to, for all $g \in \text{VIABLE}(y, \alpha)$,

$$t \in \text{SAT}(g, \alpha) \implies x \in \text{SAT}(g, \alpha).$$

We have $x \in X \setminus \text{ESC}(X, \alpha)$, since the α -exit (x, y) is non-trivial. This implies $y \notin \text{SAFESTEP}(x, \alpha)$, since $y \in S \setminus X$. Thus, we can choose $v \in \text{VIABLE}(y, \alpha)$ with $x \notin \text{SAT}(v, \alpha)$. We must then also have $t \notin \text{SAT}(v, \alpha)$, since $t \in \text{SAT}(v, \alpha) \implies x \in \text{SAT}(v, \alpha)$.

We choose $z \in S$ such that (x, z) is an edge of $G(\alpha)$, which is possible by the fact that $\alpha \in \Omega$. Notice that $x \in X \setminus \{t\}$. Indeed, we have $x \neq t$, since $x \in X \setminus \text{ESC}(X, \alpha)$ and $t \in \text{ESC}(X, \alpha)$. It follows that $z \in Y \cup \{t\} \subseteq X$ by the properties of the set Y .

We also choose $h \in \text{VIABLE}(z, \alpha)$, which is possible by lemma 3.6.

Now, in case $x = u$, we define $g := \langle (x, z), h \rangle$. In case $x \in X \setminus \{u\}$, we actually have $x \in X \setminus \{u, t\} \subseteq Y$, and we can choose a path q in $G(\alpha)$ from u to x with $t \notin N(p)$. We then define $g := \langle q, (x, z), h \rangle$. We claim that $g \in \text{ADMISS}(t, u, \alpha)$.

By lemma 3.5, we have $g \in \text{VIABLE}(u, \alpha)$. Notice that the first occurrence of x in this plan is before the first occurrence of t . Moreover, we have $y \in A(x)$ and a plan $v \in \text{VIABLE}(y, \alpha)$ with $t, x \notin \text{SAT}(v, \alpha)$. Also notice that $\text{FIRST}(v) = y$ and the state z , which follows the first occurrence of x in g , are different states, since $y \in S \setminus X$ and $z \in X$. This demonstrates that g satisfies condition *AD-iii* of admissibility, hence $g \in \text{ADMISS}(t, u, \alpha)$. \square

Corollary 3.8. *Let $\alpha \in \Omega$ and let $t \in N$. Then $\alpha_t \leq \delta(t, \alpha) < \infty$.*

Proof. We have $\text{SAFESTEP}(t, \alpha) \neq \emptyset$, by the fact that $\alpha \in \Omega$. We can then choose $u \in \text{SAFESTEP}(t, \alpha)$, and we obtain $\alpha_t \leq \beta(t, u, \alpha) \leq \delta(t, \alpha)$. By lemma 3.7, the minimization $\beta(t, u, \alpha) = \min\{\phi_t(v) \mid v \in \text{ADMISS}(t, u, \alpha)\}$ is done over a non-empty set, hence $\beta(t, u, \alpha)$ is finite for all $u \in A(t)$, which demonstrates that $\delta(t, \alpha) < \infty$. \square

Lemma 3.9. Let $\alpha \in \Omega$ and $t \in N$. Let δ be the vector that results from α when coordinate α_t is replaced with $\delta(t, \alpha)$. Then

- (i) $\text{VIABLE}(u, \delta) \subseteq \text{VIABLE}(u, \alpha)$ for all $u \in N$,
- (ii) $\text{SAFESTEP}(u, \delta) \supseteq \text{SAFESTEP}(u, \alpha)$ for all $u \in N \setminus \{t\}$,
- (iii) $B(t, \alpha) \subseteq \text{SAFESTEP}(t, \delta) \subseteq \text{SAFESTEP}(t, \alpha)$,
- (iv) $\text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \delta) \setminus \{t\}$ for all $X \subseteq N$,
- (v) $\text{ESC}(X, \alpha) \subseteq \text{ESC}(X, \delta)$ for all $X \subseteq N$ with $t \notin X$ or $t \in \text{ESC}(X, \delta)$.

Proof. Proof of (i): By corollary 3.8 we have $\delta \geq \alpha$. Then, if a plan is δ -viable, it is obviously also α -viable.

Proof of (ii): Let $u \in N \setminus \{t\}$, and let $x \in \text{SAFESTEP}(u, \alpha)$. Choose an arbitrary plan $g \in \text{VIABLE}(x, \delta)$. Then $g \in \text{VIABLE}(x, \alpha)$ by (i), and since $x \in \text{SAFESTEP}(u, \alpha)$, it follows that $u \in \text{SAT}(g, \alpha)$. Since $u \neq t$, we have $\alpha_u = \delta_u$, hence $u \in \text{SAT}(g, \delta)$. It follows that $x \in \text{SAFESTEP}(u, \delta)$.

Proof of (iii), part 1: Let $u \in B(t, \alpha)$. Choose $g \in \text{VIABLE}(u, \delta)$. To prove that $u \in \text{SAFESTEP}(t, \delta)$, it suffices to show that $t \in \text{SAT}(g, \delta)$. Case 1: If $t \notin N(g)$, then $g \in \text{ADMISS}(t, u, \alpha)$, since g is α -viable by lemma 3.9-(i) and satisfies condition *AD-i*. Then $\phi_t(g) \geq \beta(t, u, \alpha) = \delta(t, \alpha)$, where the inequality follows by the definition of $\beta(t, u, \alpha)$, and the equality follows by the choice of $u \in B(t, \alpha)$. So indeed, $t \in \text{SAT}(g, \delta)$. Case 2: If $t \in N(g)$, then $t \in \text{SAT}(g, \delta)$ follows immediately from the definition of a δ -viable plan.

Proof of (iii), part 2: Now let $u \in \text{SAFESTEP}(t, \delta)$. Choose $g \in \text{VIABLE}(u, \alpha)$. We need to show that $t \in \text{SAT}(g, \alpha)$. Case 1: If $t \notin N(g)$, then $g \in \text{VIABLE}(u, \delta)$, since $\alpha_x = \delta_x$ for all $x \in N(g)$. By the fact that $u \in \text{SAFESTEP}(t, \delta)$, it follows that $t \in \text{SAT}(g, \delta)$. Then also $t \in \text{SAT}(g, \alpha)$, since $\alpha_t \leq \delta_t$. Case 2: If $t \in N(g)$, then $t \in \text{SAT}(g, \alpha)$ is trivial.

Proof of (iv): This follows immediately from (ii).

Proof of (v): If $t \notin X$, then

$$\text{ESC}(X, \alpha) = \text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \delta) \setminus \{t\} = \text{ESC}(X, \delta).$$

If $t \in \text{ESC}(X, \delta)$, then

$$\text{ESC}(X, \alpha) \subseteq \{t\} \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \{t\} \cup (\text{ESC}(X, \delta) \setminus \{t\}) = \text{ESC}(X, \delta).$$

□

Corollary 3.10. Let $\alpha \in \Omega$, let $t \in N$, and let δ be the vector that results from α when coordinate α_t is replaced with $\delta(t, \alpha)$. Then $\text{SAFESTEP}(x, \delta) \neq \emptyset$ for all $x \in N$.

Proof. For all $x \in N \setminus \{t\}$, we have $\text{SAFESTEP}(x, \delta) \supseteq \text{SAFESTEP}(x, \alpha) \neq \emptyset$, where the inclusion follows from lemma 3.9-(ii), and the inequality follows from the fact that $\alpha \in \Omega$. We also have $\text{SAFESTEP}(t, \delta) \supseteq B(t, \alpha) \neq \emptyset$, where the inclusion follows from lemma 3.9-(iii). \square

4 Stable vectors and their properties

In the previous section we showed that for all $\alpha \in \Omega$, the updated vector, say $\delta(\alpha)$, is finite and satisfies $\delta(\alpha) \geq \alpha$. If we could also prove $\delta(\alpha) \in \Omega$ for all $\alpha \in \Omega$, then it would be an easy corollary to establish a 'fixed point' in Ω , i.e. the existence of a vector $\alpha^* \in \Omega$ with the property $\delta(t, \alpha^*) = \alpha_t^*$ for all $t \in N$.

We begin this section with an example of $\alpha \in \Omega$ and $\delta(\alpha) \notin \Omega$. The example demonstrates that, if we initiate the updating process with a vector in Ω , the process may terminate with a vector that is not finite.

This 'negative result' will motivate the rather intricate definition of the set Ω^* of stable vectors, later in this section. The set Ω^* will be a subset of Ω , so that all results derived in section 3 will also hold for all $\alpha \in \Omega^*$. Most importantly however, we will be able to prove that $\delta(\alpha) \in \Omega^*$ for all $\alpha \in \Omega^*$.




Figure 3

Example 4.1. For the game depicted in figure 3, let us set $\alpha_1 = 2$ and $\alpha_2 = \alpha_3 = \alpha_4 = \alpha_5 = 0$. We claim that α is semi-stable: It is easy to verify that the condition $\text{SAFESTEP}(x, \alpha) \neq \emptyset$ for all $x \in N$ is satisfied. Moreover, the collection $\mathcal{X}(\alpha)$ consists of only the set $\{1, 2, 3\}$, and one can verify that the edge $(3, 4)$ is a non-trivial α -exit from $\{1, 2, 3\}$.

Corollary 3.8 now predicts that any update on α will be finite. However, the vector that results after the update of the state controlled by player 4 is not semi-stable. The updated vector, say $\tilde{\alpha}$, is given by $\tilde{\alpha}_1 = \tilde{\alpha}_4 = 2$ and $\tilde{\alpha}_2 = \tilde{\alpha}_3 = \tilde{\alpha}_5 = 0$. Observe now that $\{1, 2, 3\} \in \mathcal{X}(\tilde{\alpha})$ and that $3 \in \text{ESC}(\{1, 2, 3\}, \tilde{\alpha})$. Thus, the edge $(3, 4)$ is a trivial $\tilde{\alpha}$ -exit from $\{1, 2, 3\}$. As there is no other serious candidate for a non-trivial $\tilde{\alpha}$ -exit from $\{1, 2, 3\}$, we conclude that $\tilde{\alpha}$ is not semi-stable.

Any update on $\tilde{\alpha}$ is still finite, even though corollary 3.8 does not apply. However, two consecutive updates of the vector $\tilde{\alpha}$ will result in a vector that is not finite anymore. If we first update the state controlled by 3, then its new value will be $\beta(3, 1, \tilde{\alpha}) = 1$: indeed, any $\tilde{\alpha}$ -viable plan with initial state 1 and with a reward lower than 1 for player 3 must absorb at 5, which is not $(3, 1, \tilde{\alpha})$ -admissible. If we subsequently update the state controlled by 2, we see that there are no $(2, 3)$ -admissible plans anymore, hence the value of 2 becomes infinite. The reader may wish to verify that a different order of updates does not solve the problem. \diamond

Let $\alpha \in \Omega$, $X \subseteq N$ and $Z \subseteq X$. We say that an edge (x, y) is an (α, Z) -exit from X if $x \in X$ and $y \in S \setminus X$, and if, for all $v \in \text{VIABLE}(y, \alpha)$,

$$Z \cup \text{ESC}(X, \alpha) \subseteq \text{SAT}(v, \alpha) \implies x \in \text{SAT}(v, \alpha).$$

We say that a sequence of edges $\mathfrak{e} = (x_i, y_i)_{i=1}^k$ is an α -exit sequence from X if, for all $i \in \{1, \dots, k\}$, the edge (x_i, y_i) is an $(\alpha, \{x_1, \dots, x_{i-1}\})$ -exit from X . For technical reasons, we allow $k = 0$, i.e. the empty sequence will also be called an α -exit sequence from X . We say that the α -exit sequence from X is *trivial* if $x_i \in \text{ESC}(X, \alpha)$ for all $i \in \{1, \dots, k\}$. We say that it is *non-trivial* if the sequence is non-empty and if there exists $i \in \{1, \dots, k\}$ such that $x_i \in X \setminus \text{ESC}(X, \alpha)$. We say that the α -exit sequence from X is *positive* if the sequence is non-empty and if there exists $i \in \{1, \dots, k\}$ such that $x_i \in \text{POS}(X, \alpha)$.

The following lemma states some elementary facts about exit sequences.

Lemma 4.1. *Let $X \subseteq N$ and let $\alpha \in \Omega$.*

- (i) *If \mathfrak{e} is a non-empty α -exit sequence from X , then its first edge is an α -exit from X .*
- (ii) *If \mathfrak{e} is an α -exit sequence from X , and if (x, y) is an edge of \mathfrak{e} with $x \in \text{ESC}(X, \alpha)$, then the sequence $\bar{\mathfrak{e}}$ obtained from \mathfrak{e} by deleting the edge (x, y) is also an α -exit sequence from X . Moreover, if \mathfrak{e} is a positive α -exit sequence from X and if $X \in \mathcal{X}(\alpha)$, then $\bar{\mathfrak{e}}$ is also a positive α -exit sequence from X .*
- (iii) *If a non-trivial α -exit sequence from X exists, then a non-trivial α -exit from X exists.*
- (iv) *If \mathfrak{e} and \mathfrak{f} are both α -exit sequences from X , then the concatenation of these two sequences, denoted by $(\mathfrak{e}, \mathfrak{f})$, is also an α -exit sequence from X . If moreover, \mathfrak{e} or \mathfrak{f} is a positive α -exit sequence from X , then also $(\mathfrak{e}, \mathfrak{f})$ is a positive α -exit sequence from X .*
- (v) *If \mathfrak{e} is a positive α -exit sequence from X and if $X \in \mathcal{X}(\alpha)$, then \mathfrak{e} is non-trivial.*

Proof. Proof of (i): If \mathfrak{e} is a non-empty α -exit sequence from X , then its first edge is by definition an (α, \emptyset) -exit from X . That is also the definition of an α -exit from X .

Proof of (ii): Let $\mathfrak{e} = (x_i, y_i)_{i=1}^k$ be an α -exit sequence from X . Suppose $h \in \{1, \dots, k\}$ is such that $x_h \in \text{ESC}(X, \alpha)$. Then we have, for all $j \in \{1, \dots, k\}$, that

$$\{x_1, \dots, x_{j-1}\} \cup \text{ESC}(X, \alpha) = \{x_1, \dots, x_{j-1}\} \setminus \{x_h\} \cup \text{ESC}(X, \alpha).$$

Now, by the definition of an (α, Y) -exit from X and by the fact that (x_j, y_j) is an $(\alpha, \{x_1, \dots, x_{j-1}\})$ -exit from X , it follows that edge (x_j, y_j) is an $(\alpha, \{x_1, \dots, x_{j-1}\} \setminus \{x_h\})$ -exit from X , for all $j \in \{1, \dots, k\}$; thus also for all $j \neq h$. This means that the sequence $\bar{\mathfrak{e}}$, obtained from \mathfrak{e} by deleting edge (x_h, y_h) from it, is also an α -exit sequence from X .

Moreover, if $X \in \mathcal{X}(\alpha)$, then $\text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha) = \emptyset$. So in this case we have $x_h \notin \text{POS}(X, \alpha)$, and it follows that the sequence $\bar{\mathfrak{e}}$ is positive if \mathfrak{e} is positive.

Proof of (iii): Let $\mathfrak{e} = (x_i, y_i)_{i=1}^k$ be a non-trivial α -exit sequence from X . Let h be the smallest index with $x_h \in X \setminus \text{ESC}(X, \alpha)$. Denote by $\bar{\mathfrak{e}}$ the edge-sequence obtained from \mathfrak{e} by deleting all edges (x_i, y_i) from \mathfrak{e} with $i < h$. Then the edge-sequence $\bar{\mathfrak{e}}$ is an α -exit sequence from X , by (ii). The first edge of $\bar{\mathfrak{e}}$ (i.e. (x_h, y_h)) is an α -exit from X , by (i).

Proof of (iv): Let $\mathfrak{e} = (x_i, y_i)_{i=1}^k$ and $\mathfrak{f} = (x_i, y_i)_{i=k+1}^\ell$ be two α -exit sequences from X . We need to prove that $(x_i, y_i)_{i=1}^\ell$ is an α -exit sequence from X . To see this, let $i \in \{1, \dots, \ell\}$ and let $v \in \text{VIABLE}(y_i, \alpha)$ be such that $\{x_1, \dots, x_{i-1}\} \cup \text{ESC}(X, \alpha) \subseteq \text{SAT}(v, \alpha)$. If $i \leq k$, we use the fact that (x_i, y_i) is an $(\alpha, \{x_1, \dots, x_{i-1}\})$ -exit from X to deduce that $x_i \in \text{SAT}(v, \alpha)$. If $i > k$, we use the fact that (x_i, y_i) is an $(\alpha, \{x_{k+1}, \dots, x_{i-1}\})$ -exit from X to deduce that $x_i \in \text{SAT}(v, \alpha)$.

If moreover, one of the sequences \mathfrak{e} or \mathfrak{f} is a positive α -exit sequence from X , then one of these sequence contains an edge (x, y) with $x \in \text{POS}(X, \alpha)$. Obviously, $(\mathfrak{e}, \mathfrak{f})$ also contains the edge (x, y) , and is therefore a positive α -exit sequence from X .

Proof of (v): Let $(x_i, y_i)_{i=1}^k$ be a positive α -exit sequence from $X \in \mathcal{X}(\alpha)$. Thus, there exists $i \in \{1, \dots, k\}$ with $x_i \in \text{POS}(X, \alpha)$. Then $x_i \notin \text{ESC}(X, \alpha)$, since $\text{POS}(X, \alpha) \cap \text{ESC}(X, \alpha) = \emptyset$ by the fact that $X \in \mathcal{X}(\alpha) \subseteq \mathcal{E}(\alpha)$. \square

We now say that a vector $\alpha \in \mathbb{R}^N$ is *stable* if $\text{SAFESTEP}(x, \alpha) \neq \emptyset$ for all $x \in N$ and if a positive α -exit sequence from X exists for all $X \in \mathcal{X}(\alpha)$. We denote the set of stable vectors in \mathbb{R}^N by Ω^* .

Lemma 4.2. *We have $\rho \in \Omega^* \subseteq \Omega$, where, ρ is the vector defined by $\rho_t = r_t(t)$ for all $t \in N$.*

Proof. For the vector ρ , we demonstrated in the proof of lemma 3.1 that $\text{SAFESTEP}(t, \rho) \neq \emptyset$ for all $t \in N$. We also demonstrated that $\mathcal{X}(\rho) = \emptyset$, so trivially, a positive ρ -exit sequence from X exists for all $X \in \mathcal{X}(\rho)$. Thus, $\rho \in \Omega^*$.

Since for $X \in \mathcal{X}(\alpha)$, a positive α -exit sequence from X is non-trivial (by lemma 4.1-(v)), and since the existence of a non-trivial α -exit sequence from X implies the existence of a non-trivial α -exit from X (by lemma 4.1-(iii)), it follows that any stable vector is also semi-stable. Thus,

$\Omega^* \subseteq \Omega$. □

Let us fix $\alpha \in \Omega^*$ and $t \in N$. Let further δ denote the update of α , where α_t is replaced by $\delta(t, \alpha)$. It follows from lemma 4.2 and corollary 3.10 that δ satisfies the condition $\text{SAFESTEP}(x, \delta) \neq \emptyset$ for all $x \in N$. In this section, we will demonstrate that δ also has the property that a positive δ -exit sequence from X exists, for all $X \in \mathcal{X}(\delta)$. Let us partition the set $\mathcal{X}(\delta)$ into two subsets.

$$\begin{aligned}\mathcal{V}(t, \delta) &= \{X \in \mathcal{X}(\delta) \mid t \notin X \text{ or } t \in \text{ESC}(X, \delta)\}, \\ \mathcal{W}(t, \delta) &= \{X \in \mathcal{X}(\delta) \mid t \in X \setminus \text{ESC}(X, \delta)\}.\end{aligned}$$

We will first deal with the set $\mathcal{V}(t, \delta)$. Since $\alpha \in \Omega^*$, the following lemma implies the existence of a positive α -exit sequence from X , for all $X \in \mathcal{V}(t, \delta)$.

Lemma 4.3. *We have $\mathcal{V}(t, \delta) \subseteq \mathcal{X}(\alpha)$.*

Proof. Let $X \in \mathcal{V}(t, \delta)$. We need to prove that $X \in \mathcal{X}(\alpha) = \mathcal{P}(\alpha) \cap \mathcal{E}(\alpha) \cap \mathcal{C}$.

Proof that $X \in \mathcal{P}(\alpha)$: We claim that $t \notin \text{POS}(X, \delta)$. This is trivial if $t \notin X$. Otherwise, we have $t \in \text{ESC}(X, \delta)$, since $X \in \mathcal{V}(t, \delta)$. We also have $\text{ESC}(X, \delta) \cap \text{POS}(X, \delta) = \emptyset$, by the fact that $X \in \mathcal{V}(t, \delta) \subseteq \mathcal{X}(\delta) \subseteq \mathcal{E}(\delta)$. Now $t \notin \text{POS}(X, \delta)$ follows from the combination of these facts.

Then $\text{POS}(X, \alpha) = \text{POS}(X, \delta) \neq \emptyset$, where the equality is by the fact that $\alpha_x = \delta_x$ for all $x \in X \setminus \{t\}$ and the fact that $\alpha_t \leq \delta_t \leq 0$.

Proof that $X \in \mathcal{E}(\alpha)$: We have $t \notin X$ or $t \in \text{ESC}(X, \delta)$ by the fact that $X \in \mathcal{V}(t, \delta)$. Then $\text{ESC}(X, \alpha) \subseteq \text{ESC}(X, \delta)$ by lemma 3.9-(v). It follows that

$$\text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha) \subseteq \text{ESC}(X, \delta) \cap \text{POS}(X, \delta) = \emptyset,$$

where the inclusion is by the fact that $\text{ESC}(X, \alpha) \subseteq \text{ESC}(X, \delta)$ and $\text{POS}(X, \alpha) = \text{POS}(X, \delta)$, and the equality is by the fact that $X \in \mathcal{V}(t, \delta) \subseteq \mathcal{X}(\delta) \subseteq \mathcal{E}(\delta)$.

Proof that $X \in \mathcal{C}$: Obviously, we have $X \in \mathcal{C}$, since $X \in \mathcal{V}(t, \delta) \subseteq \mathcal{X}(\delta) \subseteq \mathcal{C}$. □

Existence of a positive δ -exit sequence from X for all $X \in \mathcal{V}(t, \delta)$ is now implied by the following result.

Lemma 4.4. *For $X \in \mathcal{V}(t, \delta)$, every (positive) α -exit sequence from X is a (positive) δ -exit sequence from X .*

Proof. Let $X \in \mathcal{V}(t, \delta)$ and let $e = (x_i, y_i)_{i=1}^k$ be an α -exit sequence from X . Choose $j \in \{1, \dots, k\}$, let $g \in \text{VIABLE}(y_j, \delta)$ and assume that $\{x_1, \dots, x_{j-1}\} \cup \text{ESC}(X, \delta) \subseteq \text{SAT}(g, \delta)$. We will prove that $x_j \in \text{SAT}(g, \delta)$.

We have $\text{ESC}(X, \alpha) \subseteq \text{ESC}(X, \delta)$ by lemma 3.9-(v). Therefore,

$$\{x_1, \dots, x_{j-1}\} \cup \text{ESC}(X, \alpha) \subseteq \{x_1, \dots, x_{j-1}\} \cup \text{ESC}(X, \delta) \subseteq \text{SAT}(g, \delta) \subseteq \text{SAT}(g, \alpha).$$

Then $x_j \in \text{SAT}(g, \alpha)$, by the fact that (x_j, y_j) is an $(\alpha, \{x_1, \dots, x_{j-1}\})$ -exit from X , and by the fact that $g \in \text{VIABLE}(y_j, \delta) \subseteq \text{VIABLE}(y_j, \alpha)$. If $x_j \neq t$, then $x_j \in \text{SAT}(g, \delta)$ follows because $\alpha_{x_j} = \delta_{x_j}$. If $x_j = t$, then $x_j \in \text{SAT}(g, \delta)$ follows from the assumption $\{x_1, \dots, x_{j-1}\} \cup \text{ESC}(X, \delta) \subseteq \text{SAT}(g, \delta)$ and the fact that $t \in \text{ESC}(X, \delta)$.

So indeed, \mathfrak{e} is a δ -exit sequence from X . If \mathfrak{e} is a positive α -exit sequence from X , then \mathfrak{e} is also a positive δ -exit sequence from X , since there exists $h \in \{1, \dots, k\}$ with $x_h \in \text{POS}(X, \alpha) \subseteq \text{POS}(X, \delta)$. \square

Corollary 4.5. *For all $X \in \mathcal{V}(t, \delta)$, a positive δ -exit sequence from X exists.*

It remains to prove that a positive δ -exit sequence from X exists for all $X \in \mathcal{W}(t, \delta)$. For this, we will study α -exit sequences that are essentially independent of t . These are formally defined as follows.

Let $X \in \mathcal{W}(t, \delta)$ and let $\mathfrak{e} = (x_i, y_i)_{i=1}^k$ be an α -exit sequence from X . We will say that \mathfrak{e} is an α -exit sequence from X disregarding t if $t \notin \{x_1, \dots, x_k\}$, and if for all $i \in \{1, \dots, k\}$ and all $v \in \text{VIABLE}(y_i, \alpha)$,

$$\{x_1, \dots, x_{i-1}\} \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha) \implies x_i \in \text{SAT}(v, \alpha).$$

Lemma 4.6. *Let $X \in \mathcal{W}(t, \delta)$. Then every (positive) α -exit sequence from X disregarding t is a (positive) δ -exit sequence from X .*

Proof. Let $\mathfrak{e} = (x_i, y_i)_{i=1}^k$ be an α -exit sequence from X disregarding t .

We prove that \mathfrak{e} is a δ -exit sequence from X . The requirement $x_i \in X$ and $y_i \in S \setminus X$ for all $i \in \{1, \dots, k\}$ is obviously satisfied, because \mathfrak{e} is an α -exit sequence from X . Now, let $i \in \{1, \dots, k\}$, let $g \in \text{VIABLE}(y_i, \delta)$ and assume that $\{x_1, \dots, x_{i-1}\} \cup \text{ESC}(X, \delta) \subseteq \text{SAT}(g, \delta)$.

We have $\text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \delta)$ by lemma 3.9-(iv). Therefore,

$$\{x_1, \dots, x_{i-1}\} \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \{x_1, \dots, x_{i-1}\} \cup \text{ESC}(X, \delta) \subseteq \text{SAT}(g, \delta) \subseteq \text{SAT}(g, \alpha).$$

Then $x_i \in \text{SAT}(g, \alpha)$ follows by the fact that \mathfrak{e} is an α -exit sequence from X disregarding t . Also $x_i \neq t$ follows from that fact. The combination $x_i \in \text{SAT}(g, \alpha)$ and $x_i \neq t$ implies $x_i \in \text{SAT}(g, \delta)$. So indeed, \mathfrak{e} is a δ -exit sequence from X .

If \mathfrak{e} is a positive α -exit sequence from X disregarding t , then \mathfrak{e} is a positive δ -exit sequence from X , since there exists $h \in \{1, \dots, k\}$ with $x_h \in \text{POS}(X, \alpha) \subseteq \text{POS}(X, \delta)$. \square

In the following, if $\mathfrak{e} = (x_i, y_i)_{i=1}^k$ is a sequence of edges, we will use the notation $x_i(\mathfrak{e}) = x_i$

and $y_i(\mathbf{e}) = y_i$ for $i \in \{1, \dots, k\}$. We will also use the notation $\mathbf{x}(\mathbf{e}) = \{x_1, \dots, x_k\}$, $\mathbf{y}(\mathbf{e}) = \{y_1, \dots, y_k\}$, and $k(\mathbf{e}) = k$. We will further use the notation \emptyset for the empty sequence.

Let $X \in \mathcal{W}(t, \delta)$ and let \mathbf{e} be an α -exit sequence from X disregarding t . It will be convenient to have terminology for an edge $e = (x, y)$ such that (\mathbf{e}, e) fails to be an α -exit sequence from X disregarding t only because $y \in X$. Say that (x, y) is an α -cap for (X, \mathbf{e}, t) if $x \in X \setminus \{t\}$ and $y \in X$ and if, for all $v \in \text{VIABLE}(y, \alpha)$, we have

$$\mathbf{x}(\mathbf{e}) \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha) \implies x \in \text{SAT}(v, \alpha).$$

We denote by $F(X, \mathbf{e}, t, \alpha)$ the set of α -caps for (X, \mathbf{e}, t) . Notice that \emptyset is an α -exit sequence from X disregarding t , so $F(X, \emptyset, t, \alpha)$ is well-defined.

Lemma 4.7. *Let $X \in \mathcal{W}(t, \delta)$. Then $F(X, \emptyset, t, \alpha) \neq \emptyset$. In particular,*

(i) *for every $x \in X \setminus (\text{ESC}(X, \alpha) \cup \{t\})$, we have*

$$\emptyset \neq \{(x, y) \mid y \in \text{SAFESTEP}(x, \alpha)\} \subseteq F(X, \emptyset, t, \alpha),$$

(ii) *and for every $x \in \text{ESC}(X, \alpha) \setminus \{t\}$, we have*

$$\emptyset \neq \{(x, y) \mid y \in A(x) \cap X\} \subseteq F(X, \emptyset, t, \alpha).$$

Proof. Proof of (i): Let $x \in X \setminus (\text{ESC}(X, \alpha) \cup \{t\})$.

Non-emptiness of the set $\{(x, y) \mid y \in \text{SAFESTEP}(x, \alpha)\}$ follows from the fact that $\alpha \in \Omega^*$.

Choose $y \in \text{SAFESTEP}(x, \alpha)$. We need to show that $(x, y) \in F(X, \emptyset, t, \alpha)$. We trivially have $x \in X \setminus \{t\}$. The property $y \in X$ follows, since $x \in X \setminus \text{ESC}(X, \alpha)$, which is equivalent to $\text{SAFESTEP}(x, \alpha) \subseteq X$.

Now let $v \in \text{VIABLE}(y, \alpha)$ and assume that $\text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{SAT}(v, \alpha)$. The fact that $y \in \text{SAFESTEP}(x, \alpha)$ implies $x \in \text{SAT}(v, \alpha)$ by definition, which shows $(x, y) \in F(X, \emptyset, t, \alpha)$. (Here, the assumption $\text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{SAT}(v, \alpha)$ was not needed.)

Proof of (ii): Let $x \in \text{ESC}(X, \alpha) \setminus \{t\}$.

Non-emptiness of the set $\{(x, y) \mid y \in A(x) \cap X\}$ follows from the fact that $X \in \mathcal{W}(t, \delta) \subseteq \mathcal{C}$.

Choose $y \in A(x) \cap X$. Then trivially $x \in X \setminus \{t\}$ and $y \in X$.

Now let $v \in \text{VIABLE}(y, \alpha)$ and assume that $\text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{SAT}(v, \alpha)$. Here, $x \in \text{SAT}(v, \alpha)$ follows trivially from the assumption $\text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{SAT}(v, \alpha)$. \square

Lemma 4.8. *Let $X \in \mathcal{W}(t, \delta)$, and let \mathbf{e} and \mathbf{f} be two α -exit sequences from X disregarding t with $\mathbf{x}(\mathbf{e}) \subseteq \mathbf{x}(\mathbf{f})$. Then $F(X, \mathbf{e}, t, \alpha) \subseteq F(X, \mathbf{f}, t, \alpha)$.*

Proof. Let $(x, y) \in F(X, e, t, \alpha)$. Trivially, $x \in X \setminus \{t\}$ and $y \in X$. Now, let $v \in \text{VIABLE}(y, \alpha)$ and assume that $\mathfrak{x}(f) \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha)$. Then also $\mathfrak{x}(e) \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha)$, since $\mathfrak{x}(e) \subseteq \mathfrak{x}(f)$. Now, $x \in \text{SAT}(v, \alpha)$ follows from the fact that e is an α -exit sequence from X disregarding t . This demonstrates $(x, y) \in F(X, f, t, \alpha)$. \square

We remark that lemma 4.1 remains valid when every occurrence of the phrase "sequence from X " in the lemma is replaced by the phrase "sequence from X disregarding t ". The proof of this is completely analogous to the proof of lemma 4.1. In particular, the concatenation of two α -exit sequences from X disregarding t is an α -exit sequence from X disregarding t . Since there exists at least one α -exit sequence from X disregarding t for every $X \in \mathcal{W}(t, \delta)$ (the empty sequence), it is obvious then that, for every $X \in \mathcal{W}(t, \delta)$, an α -exit sequence e^* from X disregarding t exists such that $\mathfrak{x}(e) \subseteq \mathfrak{x}(e^*)$ for all α -exit sequences e from X disregarding t . Let us call e^* with this property a *maximal* α -exit sequence from X disregarding t . The following corollary now follows immediately from lemma 4.8.

Corollary 4.9. *Let $X \in \mathcal{W}(t, \delta)$ and let e^* be a maximal α -exit sequence from X disregarding t . Then $F(X, e, t, \alpha) \subseteq F(X, e^*, t, \alpha)$ for every α -exit sequence e from X disregarding t .*

For $X \in \mathcal{W}(t, \delta)$ and an α -exit sequence e from X disregarding t , let us define $H(X, e, t, \alpha)$ as the graph with vertex set X and edge set $F(X, e, t, \alpha) \cup \{(t, u) \mid u \in B(t, \alpha)\}$. Notice that this graph is well-defined. Indeed, for every edge in $F(X, e, t, \alpha)$, the endpoints are by definition in X . For every edge in $\{(t, u) \mid u \in B(t, \alpha)\}$, we have $t \in X$ by the fact that $X \in \mathcal{W}(t, \delta)$, and we have $u \in X$, since $u \in B(t, \alpha) \subseteq \text{SAFESTEP}(t, \delta) \subseteq X$, where the first inclusion is by lemma 3.9-(iii) and the second inclusion is by the fact that $t \in X \setminus \text{ESC}(X, \delta)$.

Lemma 4.10. *Let $X \in \mathcal{W}(t, \delta)$, let e be an α -exit sequence from X disregarding t , let p be a path in $H(X, e, t, \alpha)$, and let $g \in \text{VIABLE}(\text{LAST}(p), \alpha)$. Then*

$$\mathfrak{x}(e) \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(g, \alpha) \implies \langle p, g \rangle \in \text{VIABLE}(\text{FIRST}(p), \alpha).$$

Proof. Assume that $\mathfrak{x}(e) \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(g, \alpha)$. We write $p = (z_i)_{i=1}^m$, and we define $g_m = g$ and $g_i = \langle (z_i, \dots, z_m), g \rangle$ for $i \in \{1, \dots, m-1\}$.

Trivially, $g_m = g$ is α -viable. Now, let $i \in \{1, \dots, m-1\}$ and assume that $g_{i+1} \in \text{VIABLE}(z_{i+1}, \alpha)$. It follows immediately from $(z_i, z_{i+1}) \in F(X, e, t, \alpha)$ and the assumption $\mathfrak{x}(e) \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(g, \alpha) = \text{SAT}(g_{i+1}, \alpha)$ that $z_i \in \text{SAT}(g_{i+1}, \alpha)$. Thus, $\text{N}(g_i) = \{z_i\} \cup \text{N}(g_{i+1}) \subseteq \text{SAT}(g_{i+1}, \alpha) = \text{SAT}(g_i, \alpha)$. This proves that g_i is α -viable.

We proved by induction that g_i is α -viable for all $i \in \{1, \dots, m\}$. In particular, $\langle p, g \rangle = g_1$ is α -viable. \square

Let $X \in \mathcal{W}(t, \delta)$ and let e be an α -exit sequence from X disregarding t . Notice that, for every $x \in X$, there exists $y \in X$ such that (x, y) is an edge of $H(X, e, t, \alpha)$. This follows directly from

lemma 4.7 and the fact that $B(t, \alpha) \neq \emptyset$. Thus, by lemma 3.3, the graph $H(X, \mathfrak{e}, t, \alpha)$ has at least one ergodic set, and for every $x \in X$, there is a path in $H(X, \mathfrak{e}, t, \alpha)$ from x to an ergodic set of $H(X, \mathfrak{e}, t, \alpha)$.

Lemma 4.11. *Let $X \in \mathcal{W}(t, \delta)$ and let \mathfrak{e} be an α -exit sequence from X disregarding t such that $\mathfrak{x}(\mathfrak{e}) \cap \text{POS}(X, \alpha) = \emptyset$. Then $H(X, \mathfrak{e}, t, \alpha)$ has an ergodic set Y with $\text{POS}(Y, \alpha) \neq \emptyset$.*

Proof. In the proof, we denote $H(X, \mathfrak{e}, t, \alpha)$ by H .

Assume that $\text{POS}(Y, \alpha) = \emptyset$ for every ergodic set Y of H . We will demonstrate then that $\text{POS}(X, \delta) = \emptyset$. Let $x \in X$.

First consider the case that x is element of an ergodic set Y of H . Then obviously $\alpha_x \leq 0$, by the assumption that $\text{POS}(Y, \alpha) = \emptyset$. If $x \neq t$, it follows immediately that $\delta_x = \alpha_x \leq 0$. If $x = t$, we choose $u \in B(t, \alpha)$. Notice that $u \in Y$, since Y is an ergodic set of H and (t, u) is an edge of H . We can therefore choose a non-absorbing plan g with $\text{N}(g) \subseteq Y$ (possible by the properties of an ergodic set) and with $\text{FIRST}(g) = u$ (possible, since $u \in Y$). Plan g is α -viable by the assumption that $\text{POS}(Y, \alpha) = \emptyset$. Then g is also (t, u, α) -admissible as it satisfies AD -i. It follows that $\delta_t = \delta(t, \alpha) = \beta(t, u, \alpha) \leq \phi_t(g) = 0$.

Now consider the case that x is not an element of any ergodic set of H . Then choose an ergodic set Y of H such that there exists a path p in H from x to an element of Y . We also choose a non-absorbing plan g with $\text{N}(g) \subseteq Y$ and with $\text{FIRST}(g) = \text{LAST}(p)$. Plan g is α -viable by the assumption $\text{POS}(Y, \alpha) = \emptyset$. We claim that plan $\langle p, g \rangle$ is also α -viable.

Notice that $\mathfrak{x}(\mathfrak{e}) \subseteq \text{SAT}(g, \alpha)$ by the assumption $\mathfrak{x}(\mathfrak{e}) \cap \text{POS}(X, \alpha) = \emptyset$. We also have

$$(\text{ESC}(X, \alpha) \setminus \{t\}) \cap \text{POS}(X, \alpha) \subseteq \text{ESC}(X, \delta) \cap \text{POS}(X, \delta) = \emptyset,$$

where the inclusion follows by $\text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \delta)$ (lemma 3.9-(iv)) and $\text{POS}(X, \alpha) \subseteq \text{POS}(X, \delta)$, and the equality follows by the fact that $X \in \mathcal{W}(t, \delta) \subseteq \mathcal{X}(\delta) \subseteq \mathcal{E}(\delta)$. Thus, $\text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{SAT}(g, \alpha)$. Then $\langle p, g \rangle$ is α -viable by lemma 4.10.

It follows that $\alpha_x \leq \phi_x(\langle p, g \rangle) = 0$. If $x \neq t$, we obtain immediately $\delta_x = \alpha_x \leq 0$. If $x = t$, consider the plan $\langle p', g \rangle$, where p' is the part of p that starts at the second state of p . Clearly, the plan $\langle p', g \rangle$ is α -viable. Then $\langle p', g \rangle$ is also $(t, \text{FIRST}(p'), \alpha)$ -admissible as it satisfies AD -i. Further notice that $\text{FIRST}(p') \in B(t, \alpha)$. Thus, $\delta_t = \delta(t, \alpha) = \beta(t, \text{FIRST}(p'), \alpha) \leq \phi_t(\langle p', g \rangle) = 0$. So, we have indeed $\text{POS}(X, \delta) = \emptyset$. This contradicts $X \in \mathcal{W}(t, \delta) \subseteq \mathcal{X}(\delta) \subseteq \mathcal{P}(\delta)$. \square

Lemma 4.12. *Let $X \in \mathcal{W}(t, \delta)$, let \mathfrak{e} be an α -exit sequence from X disregarding t , and let Y be an ergodic set of $H(X, \mathfrak{e}, t, \alpha)$. Then*

$$\text{ESC}(Y, \alpha) \setminus \{t\} = (Y \cap \text{ESC}(X, \alpha)) \setminus \{t\}.$$

Proof. In the proof, we denote $H(X, \mathfrak{e}, t, \alpha)$ by H .

To see that $\text{ESC}(Y, \alpha) \setminus \{t\} \subseteq (Y \cap \text{ESC}(X, \alpha)) \setminus \{t\}$, let $x \in \text{ESC}(Y, \alpha) \setminus \{t\}$. By definition of the set $\text{ESC}(Y, \alpha)$, we can choose $y \in S \setminus Y$, such that $y \in \text{SAFESTEP}(x, \alpha)$. Now suppose $y \in X$. Then, since $x \neq t$, (x, y) is an edge of H , which contradicts that Y is an ergodic set of H . Therefore, $x \in \text{ESC}(X, \alpha)$, hence $x \in (Y \cap \text{ESC}(X, \alpha)) \setminus \{t\}$.

To see that $\text{ESC}(Y, \alpha) \setminus \{t\} \supseteq (Y \cap \text{ESC}(X, \alpha)) \setminus \{t\}$, let $x \in (Y \cap \text{ESC}(X, \alpha)) \setminus \{t\}$. By definition of the set $\text{ESC}(X, \alpha)$, we can choose $y \in S \setminus X$, such that $y \in \text{SAFESTEP}(x, \alpha)$. Then obviously $y \in S \setminus Y$, since $Y \subseteq X$. It follows that $x \in \text{ESC}(Y, \alpha)$, since $x \in Y$. Subsequently, it follows that $x \in \text{ESC}(Y, \alpha) \setminus \{t\}$, since $x \neq t$. \square

Lemma 4.13. *Let $X \in \mathcal{W}(t, \delta)$ and let e^* be a maximal α -exit sequence from X disregarding t . Then one of the following holds.*

(i) e^* is a positive α -exit sequence from X ,

(ii) for all $v \in \text{VIABLE}(t, \alpha)$,

$$\mathfrak{x}(e^*) \cup \text{ESC}(X, \alpha) \subseteq \text{SAT}(v, \alpha) \implies t \in \text{SAT}(v, \delta).$$

Proof. If e^* is a positive α -exit sequence from X , we are done. We assume from here that this is not the case. Then lemma 4.11 applies, and we can choose an ergodic set Y of $H = H(X, e^*, t, \alpha)$ with $\text{POS}(Y, \alpha) \neq \emptyset$.

A: We will first prove that $t \in Y$. Suppose that $t \notin Y$.

A1: We claim that $Y \in \mathcal{X}(\alpha)$. We trivially have $Y \in \mathcal{P}(\alpha)$, and by the properties of an ergodic set, we have $Y \in \mathcal{C}$. It remains to prove that $Y \in \mathcal{E}(\alpha)$.

We obviously have $\text{POS}(Y, \alpha) \subseteq \text{POS}(X, \delta)$. We also have

$$\text{ESC}(Y, \alpha) = \text{ESC}(Y, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \delta) \setminus \{t\} \subseteq \text{ESC}(X, \delta).$$

Here, the equality is because $t \notin Y$, the first inclusion is by lemma 4.12, the second inclusion follows by 3.9-(iv), and the third inclusion is trivial. We now conclude that

$$\text{ESC}(Y, \alpha) \cap \text{POS}(Y, \alpha) \subseteq \text{ESC}(X, \delta) \cap \text{POS}(X, \delta) = \emptyset,$$

where the equality is because $X \in \mathcal{X}(\delta) \subseteq \mathcal{E}(\delta)$. This proves that $Y \in \mathcal{E}(\alpha)$.

A2: Since $Y \in \mathcal{X}(\alpha)$ and since $\alpha \in \Omega^*$, we can choose a positive α -exit sequence e from Y . We claim that e is an α -exit sequence from X disregarding t .

We have $\mathfrak{x}(e) \subseteq Y \subseteq X \setminus \{t\}$, where the first inclusion follows by the fact that e is an α -exit sequence from Y , and the second inclusion is by the fact that $Y \subseteq X$ and $t \notin Y$.

By definition of an α -exit sequence from Y , the sequence e satisfies, for all $i \in \{1, \dots, k(e)\}$,

and all $v \in \text{VIABLE}(y_i(\mathbf{e}), \alpha)$,

$$\{x_1(\mathbf{e}), \dots, x_{i-1}(\mathbf{e})\} \cup \text{ESC}(Y, \alpha) \subseteq \text{SAT}(v, \alpha) \implies x_i(\mathbf{e}) \in \text{SAT}(v, \alpha).$$

We have $\text{ESC}(Y, \alpha) = \text{ESC}(Y, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \alpha) \setminus \{t\}$, where the equality is by the fact that $t \notin Y$, and the inclusion is by lemma 4.12. It follows that, for all $i \in \{1, \dots, k(\mathbf{e})\}$ and all $v \in \text{VIABLE}(y_i(\mathbf{e}), \alpha)$,

$$\{x_1(\mathbf{e}), \dots, x_{i-1}(\mathbf{e})\} \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha) \implies x_i(\mathbf{e}) \in \text{SAT}(v, \alpha).$$

It remains to prove that $y(\mathbf{e}) \subseteq S \setminus X$. Assume that this is not true. Since \mathbf{e} is an α -exit sequence from Y , we have $y(\mathbf{e}) \subseteq S \setminus Y = (X \setminus Y) \cup (S \setminus X)$. Then, by our assumption $y(\mathbf{e}) \not\subseteq S \setminus X$, there must be $h \in \{1, \dots, k(\mathbf{e})\}$ with the property $y_h(\mathbf{e}) \in X \setminus Y \subseteq X$. Choose the smallest h with this property. Let $f = (x_h(\mathbf{e}), y_h(\mathbf{e}))$, and let $\mathbf{f} = (x_j(\mathbf{e}), y_j(\mathbf{e}))_{j=1}^{h-1}$.

Observe then that \mathbf{f} is an α -exit sequence from X disregarding t and that f is an α -cap for (X, \mathbf{f}, t) . Thus, $f \in F(X, \mathbf{f}, t, \alpha) \subseteq F(X, \mathbf{e}^*, t, \alpha)$, where the inclusion is by lemma 4.9. We also have $\text{FIRST}(f) \in Y$ and $\text{LAST}(f) \in S \setminus Y$, as f is an edge from an α -exit sequence from Y . The set Y is an ergodic set of the graph H , so f is *not* an edge of H . Thus, $f \notin F(X, \mathbf{e}^*, t, \alpha)$. Contradiction.

So indeed, \mathbf{e} is an α -exit sequence from X disregarding t . Then it is even a positive α -exit sequence from X disregarding t , since it is a positive α -exit sequence from Y . It also follows that \mathbf{e}^* is a positive α -exit sequence from X disregarding t , by the maximality of \mathbf{e}^* . Contradiction.

B: So we have indeed $t \in Y$. We will prove by contradiction that condition (ii) of the lemma holds. Assume therefore that $g \in \text{VIABLE}(t, \alpha)$ exists with $x(\mathbf{e}^*) \cup \text{ESC}(X, \alpha) \subseteq \text{SAT}(g, \alpha)$ and $t \notin \text{SAT}(g, \delta)$.

B-1: Proof that $t \notin \text{POS}(Y, \alpha)$.

Choose $u \in B(t, \alpha)$. Since $t \in Y$ and since (t, u) is an edge of the graph H , it follows that $u \in Y$, by the properties of an ergodic set. Also by the properties of an ergodic set, a path p in H from u to t exists with $N(p) \subseteq Y$. Notice that $\langle p, g \rangle \in \text{VIABLE}(u, \alpha)$, by lemma 4.10.

Notice that $t \in \text{SAT}(v, \delta)$ for all $v \in \text{ADMISS}(t, u, \alpha)$, since we chose $u \in B(t, \alpha)$. Apparently, $\langle p, g \rangle \notin \text{ADMISS}(t, u, \alpha)$, as we have $t \notin \text{SAT}(g, \delta) = \text{SAT}(\langle p, g \rangle, \delta)$.

Then, plan $\langle p, g \rangle$ violates all three conditions *AD-i*, *AD-ii*, and *AD-iii*. Now, from the fact that $\langle p, g \rangle$ violates *AD-ii*, we deduce that $t \notin \text{POS}(Y, \alpha)$.

B-2: Proof that

$$\forall x \in Y \setminus \{t\}, \forall y \in A(x) \setminus Y, \forall v \in \text{VIABLE}(y, \alpha) : t \in \text{SAT}(v, \alpha) \vee x \in \text{SAT}(v, \alpha).$$

We further exploit the properties of plan g and path p , chosen in **B-1** to prove that $t \notin \text{POS}(Y, \alpha)$.

Choose $x \in N(p) \setminus \{t\}$. Now, from the fact that $\langle p, g \rangle$ violates *AD-iii*, we deduce that

$$\forall y \in A(x) \setminus Y, \forall v \in \text{VIABLE}(y, \alpha) : t \in \text{SAT}(v, \alpha) \vee x \in \text{SAT}(v, \alpha), \quad (1)$$

as this formally negates the existence of a 'threat' plan for player x that starts at an element of $A(x) \setminus Y$. (Notice that the elements of $A(x) \setminus Y$ do not coincide with the follower of x on path p , since $N(p) \subseteq Y$.) We need to prove that equation (1) not only holds for $x \in N(p) \setminus \{t\}$, but for all $x \in Y \setminus \{t\}$.

Choose $x \in Y \setminus \{t\}$ arbitrarily. By the properties of an ergodic set and the fact that $t \in Y$, a path q in H exists from t to x with $N(q) \subseteq Y$. Obviously, we may require that there is only one occurrence of t on this path, at the beginning. Let \tilde{q} denote the part of q that starts at the second state. There also exists a path r in H from x to t with $N(r) \subseteq Y$, and we may require that there is only one occurrence of t on this path, at the end. Let $p' = \langle \tilde{q}, r \rangle$. Clearly, p' is a path in H from $\text{FIRST}(p')$ to t with $N(p') \subseteq Y$ and with only one occurrence of t on this path, at the end. We claim that $\text{FIRST}(p') \in B(t, \alpha)$. Indeed, this is true, because $(t, \text{FIRST}(p'))$ is an edge of the graph H and edges of the type (t, u) in H all have the property $u \in B(t, \alpha)$.

We now let p' play the role of the earlier chosen path p , and argue similarly that $\langle p', g \rangle$ is an α -viable plan, such that $\langle p', g \rangle \notin \text{ADMISS}(t, \text{FIRST}(p'), \alpha)$. From the fact that plan $\langle p', g \rangle$ violates *AD-iii*, and the fact that plan $\langle p', g \rangle$ is constructed such that the first occurrence of x is before t , it follows that equation (1) holds indeed.

B-3: Proof that $Y \in \mathcal{X}(\alpha)$.

We trivially have $Y \in \mathcal{P}(\alpha)$. We also have $Y \in \mathcal{C}$ by the properties of an ergodic set. Since $t \notin \text{POS}(Y, \alpha)$, we may write

$$\text{ESC}(Y, \alpha) \cap \text{POS}(Y, \alpha) = (\text{ESC}(Y, \alpha) \setminus \{t\}) \cap \text{POS}(Y, \alpha).$$

We have

$$\text{ESC}(Y, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \delta) \setminus \{t\} \subseteq \text{ESC}(X, \delta),$$

where the first inclusion is by lemma 4.12 and the second inclusion lemma 3.9-(iv). We also have $\text{POS}(Y, \alpha) \subseteq \text{POS}(X, \delta)$, by the fact that $Y \subseteq X$ and $\alpha \leq \delta$. Therefore,

$$\text{ESC}(Y, \alpha) \cap \text{POS}(Y, \alpha) = (\text{ESC}(Y, \alpha) \setminus \{t\}) \cap \text{POS}(Y, \alpha) \subseteq \text{ESC}(X, \delta) \cap \text{POS}(X, \delta) = \emptyset,$$

where the last equality is by the fact that $X \in \mathcal{X}(\delta) \subseteq \mathcal{E}(\delta)$. This proves that $Y \in \mathcal{E}(\alpha)$, hence $Y \in \mathcal{P}(\alpha) \cap \mathcal{E}(\alpha) \cap \mathcal{C} = \mathcal{X}(\alpha)$.

B-4: Existence of a positive α -exit sequence from Y disregarding t .

Since $Y \in \mathcal{X}(\alpha)$ and since $\alpha \in \Omega^*$, we can choose a positive α -exit sequence \bar{e} from Y . Let \bar{e} denote the sequence that results from \bar{e} by deleting all edges $(x_j(\bar{e}), y_j(\bar{e}))$ with the property

$x_j(\bar{e}) = t$. We claim that e is a positive α -exit sequence from Y disregarding t .

Since \bar{e} is an α -exit sequence from Y , we have $\mathfrak{x}(\bar{e}) \subseteq Y$. Then $\mathfrak{x}(e) \subseteq Y \setminus \{t\}$ due to the construction of e . Obviously, we have $\mathfrak{y}(\bar{e}) \subseteq S \setminus Y$, as \bar{e} is an α -exit sequence from Y .

Let $j \in \{1, \dots, k(e)\}$, let $v \in \text{VIABLE}(y_j(e))$, and assume that

$$\{x_1(e), \dots, x_{j-1}(e)\} \cup (\text{ESC}(Y, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha).$$

Notice that the result of **B-2** applies here, since $x_j(e) \in Y \setminus \{t\}$ and $y_j(e) \in A(x) \setminus Y$. Thus, we have $t \in \text{SAT}(v, \alpha)$ or $x_j(e) \in \text{SAT}(v, \alpha)$. In fact, we have $x_j(e) \in \text{SAT}(v, \alpha)$. Indeed, in the case $t \in \text{SAT}(v, \alpha)$, we have

$$\{x_1(\bar{e}), \dots, x_{i-1}(\bar{e})\} \cup \text{ESC}(Y, \alpha) \subseteq \text{SAT}(v, \alpha),$$

where $i \in \{1, \dots, k(\bar{e})\}$ is such that $x_i(\bar{e}) = x_j(e)$. It then follows that $x_j(e) = x_i(\bar{e}) \in \text{SAT}(v, \alpha)$, by the fact that $(x_i(\bar{e}), y_i(\bar{e}))$ is a $(\alpha, \{x_1(\bar{e}), \dots, x_{i-1}(\bar{e})\})$ -exit from Y . Thus, for all $j \in \{1, \dots, k(e)\}$ and for all $v \in \text{VIABLE}(y_j(e))$, we have

$$\{x_1(e), \dots, x_{j-1}(e)\} \cup (\text{ESC}(Y, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha) \implies x_j(e) \in \text{SAT}(v, \alpha).$$

This proves that e is an α -exit sequence from Y disregarding t .

Since \bar{e} is a positive α -exit sequence from Y , we have $\mathfrak{x}(\bar{e}) \cap \text{POS}(X, \alpha) \neq \emptyset$. We have $t \notin \text{POS}(X, \alpha)$, by the result of **B-1**. Therefore, $\mathfrak{x}(e) \cap \text{POS}(X, \alpha) = \mathfrak{x}(\bar{e}) \cap \text{POS}(X, \alpha) \neq \emptyset$. It follows that e is a positive α -exit sequence from Y disregarding t . The sequence e will now be used to derive a contradiction.

B-5: Derivation of the contradiction.

We distinguish between two cases.

Case 1: Suppose that $\mathfrak{y}(e) \subseteq S \setminus X$. We then claim that e is a positive α -exit sequence from X disregarding t . For this, it remains to prove that, for all $j \in \{1, \dots, k(e)\}$ and for all $v \in \text{VIABLE}(y_j(e))$, we have

$$\{x_1(e), \dots, x_{j-1}(e)\} \cup (\text{ESC}(X, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha) \implies x_j(e) \in \text{SAT}(v, \alpha).$$

Indeed, this is true because (i), we have

$$\{x_1(e), \dots, x_{j-1}(e)\} \cup (\text{ESC}(Y, \alpha) \setminus \{t\}) \subseteq \text{SAT}(v, \alpha) \implies x_j(e) \in \text{SAT}(v, \alpha).$$

by the fact that e is an α -exit sequence from Y disregarding t , and because (ii), we have $\text{ESC}(Y, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \alpha) \setminus \{t\}$, by lemma 4.12.

So e is a positive α -exit sequence from X disregarding t . Then e^* is also a positive α -exit

sequence from X disregarding t , since we have $\mathfrak{x}(e^*) \cap \text{POS}(X, \alpha) \supseteq \mathfrak{x}(e) \cap \text{POS}(X, \alpha) \neq \emptyset$ by the maximality of e^* . This contradicts our assumption at the beginning of the lemma.

Case 2: Suppose that $y(e) \not\subseteq S \setminus X$. Then there must be $h \in \{1, \dots, k(e)\}$ with the property $y_h(e) \in X$. Choose the smallest h with this property. Let $f = (x_h(e), y_h(e))$, and let $\mathfrak{f} = (x_j(e), y_j(e))_{j=1}^{h-1}$.

Observe then that \mathfrak{f} is an α -exit sequence from X disregarding t and that f is an α -cap for (X, \mathfrak{f}, t) . Then $f \in F(X, \mathfrak{f}, t, \alpha) \subseteq F(X, e^*, t, \alpha)$, where the inclusion is by lemma 4.9. We also have $\text{FIRST}(f) \in Y$ and $\text{LAST}(f) \in S \setminus Y$, as f is an edge from an α -exit sequence from Y . The set Y is an ergodic set of the graph H , so f is *not* an edge of H . Thus, $f \notin F(X, e^*, t, \alpha)$. Contradiction. \square

Lemma 4.14. *For all $X \in \mathcal{W}(t, \delta)$, a positive δ -exit sequence from X exists.*

Proof. Let $X \in \mathcal{W}(t, \delta)$ and let e^* be a maximal α -exit sequence from X disregarding t . If e^* is a positive α -exit sequence from X disregarding t , then it is also a positive δ -exit sequence from X by lemma 4.6, and we are done. Assume from here that e^* is not a positive α -exit sequence from X disregarding t . Then, by lemma 4.13, we have for all $v \in \text{VIABLE}(t, \alpha)$,

$$\mathfrak{x}(e^*) \cup \text{ESC}(X, \alpha) \subseteq \text{SAT}(v, \alpha) \implies t \in \text{SAT}(v, \delta). \quad (2)$$

We distinguish between the cases $t \in \text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha)$ and $t \notin \text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha)$.

First suppose $t \in \text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha)$. Since $t \in \text{ESC}(X, \alpha)$, we can choose $s \in S \setminus X$, such that $s \in \text{SAFESTEP}(t, \alpha)$. We claim that $(e^*, (t, s))$ is a δ -exit sequence from X . For this, we only need to check that (t, s) is a $(\mathfrak{x}(e^*), \delta)$ -exit from X , since we already know that e^* is a δ -exit sequence from X , by lemma 4.6. We have indeed $t \in \text{SAT}(v, \delta)$ for every $v \in \text{VIABLE}(s, \delta)$ with $\mathfrak{x}(e^*) \cup \text{ESC}(X, \delta) \subseteq \text{SAT}(v, \delta)$, because e^* satisfies condition (ii) of lemma 4.13 and because $\text{ESC}(X, \delta) \supseteq \text{ESC}(X, \alpha) \setminus \{t\}$ by lemma 3.9-(iv). Then $(e^*, (t, s))$ is also a positive δ -exit sequence from X , since $t \in \text{POS}(X, \alpha) \subseteq \text{POS}(X, \delta)$.

Now suppose $t \notin \text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha)$. We claim that $X \in \mathcal{X}(\alpha)$.

According to lemma 4.11, the graph $H(X, e^*, t, \alpha)$ has an ergodic set Y such that $\text{POS}(Y, \alpha) \neq \emptyset$. It follows that $\text{POS}(X, \alpha) \supseteq \text{POS}(Y, \alpha) \neq \emptyset$. Thus, $X \in \mathcal{P}(\alpha)$.

We also have $X \in \mathcal{E}(\alpha)$, since

$$\text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha) = (\text{ESC}(X, \alpha) \setminus \{t\}) \cap \text{POS}(X, \alpha) \subseteq \text{ESC}(X, \delta) \cap \text{POS}(X, \delta) = \emptyset,$$

where the first equality follows, since $t \notin \text{ESC}(X, \alpha) \cap \text{POS}(X, \alpha)$.

Obviously, $X \in \mathcal{C}$, since $X \in \mathcal{X}(\delta) \subseteq \mathcal{C}$. So indeed, $X \in \mathcal{P}(\alpha) \cap \mathcal{E}(\alpha) \cap \mathcal{C} = \mathcal{X}(\alpha)$.

Now, since $\alpha \in \Omega^*$ and $X \in \mathcal{X}(\alpha)$, we can choose a positive α -exit sequence \mathfrak{f} from X . If $t \in \text{ESC}(X, \alpha)$, we require that $t \notin \mathfrak{x}(\mathfrak{f})$, which is possible by lemma 4.1-(ii). We claim that

$(\mathbf{e}^*, \mathbf{f})$ is a positive δ -exit sequence from X .

It is obvious that an edge in $(\mathbf{e}^*, \mathbf{f})$ of the form $(x_j(\mathbf{e}^*), y_j(\mathbf{e}^*))$ is a $(\delta, \{x_1(\mathbf{e}^*), \dots, x_{j-1}(\mathbf{e}^*)\})$ -exit from X , as \mathbf{e}^* is a δ -exit sequence from X , by lemma 4.6.

It is straightforward to prove that an edge in $(\mathbf{e}^*, \mathbf{f})$ of the form $(x_j(\mathbf{f}), y_j(\mathbf{f}))$ and with $x_j(\mathbf{f}) \neq t$ is a $(\delta, \{x_1(\mathbf{f}), \dots, x_{j-1}(\mathbf{f})\})$ -exit from X , using the fact that \mathbf{f} is an α -exit sequence from X . Then the edge is also a $(\delta, \mathbf{x}(\mathbf{e}^*) \cup \{x_1(\mathbf{f}), \dots, x_{j-1}(\mathbf{f})\})$ -exit from X .

It remains to prove for an edge in \mathbf{f} of the type (t, s) , say $(t, s) = (x_j(\mathbf{f}), y_j(\mathbf{f}))$, that it is a $(\delta, \mathbf{x}(\mathbf{e}^*) \cup \{x_1(\mathbf{f}), \dots, x_{j-1}(\mathbf{f})\})$ -exit from X . Let $v \in \text{VIABLE}(s, \delta)$, and assume that

$$\mathbf{x}(\mathbf{e}^*) \cup \{x_1(\mathbf{f}), \dots, x_{j-1}(\mathbf{f})\} \cup \text{ESC}(X, \delta) \subseteq \text{SAT}(v, \delta). \quad (3)$$

We need to prove that $t \in \text{SAT}(v, \delta)$. The fact that t appears in the α -exit sequence \mathbf{f} implies that $t \notin \text{ESC}(X, \alpha)$, by the choice of \mathbf{f} . Thus,

$$\text{ESC}(X, \alpha) = \text{ESC}(X, \alpha) \setminus \{t\} \subseteq \text{ESC}(X, \delta) \setminus \{t\} \subseteq \text{ESC}(X, \delta),$$

where the first inclusion is by lemma 3.9-(iv). Assumption (3) therefore implies that

$$\{x_1(\mathbf{f}), \dots, x_{j-1}(\mathbf{f})\} \cup \text{ESC}(X, \alpha) \subseteq \text{SAT}(v, \delta) \subseteq \text{SAT}(v, \alpha).$$

We then obtain that $t = x_j(\mathbf{f}) \in \text{SAT}(v, \alpha)$ from the fact that \mathbf{f} is an α -exit sequence from X and the fact that $v \in \text{VIABLE}(s, \delta) \subseteq \text{VIABLE}(s, \alpha)$.

Assumption (3) also implies

$$\mathbf{x}(\mathbf{e}^*) \cup \text{ESC}(X, \alpha) \subseteq \text{SAT}(v, \alpha) = \text{SAT}(\langle (t, s), v \rangle, \alpha).$$

Notice that the plan $\langle (t, s), v \rangle$ is α -viable, since we just proved that $t \in \text{SAT}(v, \alpha)$. Then, by equation (2), it subsequently follows that $t \in \text{SAT}(v, \delta)$.

We proved that $(\mathbf{e}^*, \mathbf{f})$ is a δ -exit sequence from X . It is a positive δ -exit sequence from X , since $\mathbf{x}((\mathbf{e}^*, \mathbf{f})) \cap \text{POS}(X, \delta) \supseteq \mathbf{x}(\mathbf{f}) \cap \text{POS}(X, \alpha) \neq \emptyset$. \square

The following result is a direct consequence of corollary 4.5 and lemma 4.14.

Corollary 4.15. $\delta \in \Omega^*$.

We have arrived at the main result of this section.

Theorem 4.16. *There exists a vector $\alpha^* \in \Omega^*$ such that $\delta(t, \alpha^*) = \alpha_t^*$ for all $t \in N$.*

Proof. We construct a sequence $(\alpha^k)_{k \in \mathbb{N}}$ with $\alpha^k \in \mathbb{R}^N$ for all $k \in \mathbb{N}$ as follows.

We set $\alpha^1 = \rho$, where ρ is the vector defined by $\rho_t = r_t(t)$ for all $t \in N$. For $k > 1$, the vector α^k is obtained from α^{k-1} by replacing one of its coordinates, say t , by the updated value $\delta(t, \alpha^{k-1})$. For the coordinate t , we choose the smallest one with the property $\delta(t, \alpha^{k-1}) \neq \alpha_t^{k-1}$ if there is such a coordinate, and we choose $t = 1$ otherwise. (Since $N = \{1, \dots, n\}$, there is a natural ordering of the elements of N , which we use here.)

Notice that this indeed defines an infinite sequence of vectors in \mathbb{R}^N (i.e. the construction of the sequence will not halt due to a coordinate that becomes infinite): we have $\alpha^1 = \rho \in \Omega^* \subseteq \mathbb{R}^N$ by lemma 4.2, and we obtain $\alpha^k \in \Omega^* \subseteq \mathbb{R}^N$ for all $k > 1$ by induction with corollary 4.15.

Further, notice that, for any $\alpha \in \Omega^*$ and $t \in N$, the number $\delta(t, \alpha)$ can only take values in the set $V_t = \{0\} \cup \{r_t(s) \mid s \in N\}$, since $\delta(t, \alpha)$ is defined as the average reward for t associated with some plan. We have $\alpha_t^1 = \rho_t = r_t(t) \in V_t$ for all $t \in N$. Thus, we obtain $\alpha_t^k \in V_t$ for all $k \in \mathbb{N}$, by induction.

If $\alpha^k \neq \alpha^{k-1}$ for $k > 1$, then the difference between these two vectors is in one coordinate. By corollary 3.8, there is an increase in that coordinate when we go from α^{k-1} to α^k . There can be at most n increases of coordinate t , for any $t \in N$, since $\alpha_t^k \in V_t$ for all $k \in \mathbb{N}$ and since V_t has cardinality $n + 1$. As there are $|N| = n$ different coordinates, the total number of increases is bounded by n^2 . It follows that the sequence $(\alpha^k)_{k \in \mathbb{N}}$ is eventually constant.

Thus, we can choose k such that $\alpha^k = \alpha^{k-1}$. This means that $\delta(t, \alpha^{k-1}) = \alpha_t^{k-1}$ for all $t \in N$, i.e. α^{k-1} satisfies the requirements of the theorem. \square

References

- [1] Flesch J., Kuipers J., Schoenmakers G. and Vrieze K. (2010), Subgame Perfection in Positive Recursive Games with Perfect Information, *Math. of OR*, **35(4)**, 742-755.
- [2] Flesch J., Kuipers J., Mashiah-Yaakovi A., Schoenmakers G., Solan E. and Vrieze K. (2010), Perfect-Information Games with Lower-Semicontinuous Payoffs, *Math. of OR*, **35(1)**, 193-207.
- [3] Kuipers J., Flesch J., Schoenmakers G. and Vrieze K. (2013), Subgame Perfection in Free Transition Games, *European Journal of Operational Research*, **228**, 201-207.
- [4] Flesch J., Kuipers J., Mashiah-Yaakovi A., Schoenmakers G., Shmaya E., Solan E. and Vrieze K. (2014), Non-existence of subgame-perfect ε -equilibrium in perfect information games with infinite horizon, to appear in: *International Journal of Game Theory*.
- [5] Laraki R., Maitra A.P. and Sudderth W. D. (2013), Two-Person Zero-Sum Stochastic Games with Semicontinuous Payoff, *Dynamic Games and Applications*, **3(2)**, 162-171.

- [6] Mertens J.F. and Neyman A. (1981), Stochastic Games, *International Journal of Game Theory*, **10(2)**, 53-66.
- [7] Mertens J.F. (1987), Repeated Games, *Proceedings of the International Congress of Mathematicians*, **1**, (Berkeley, California, 1986), 1528-1577, American Mathical Society, Providence, RI.
- [8] Purves R.A. and Sudderth W.D. (2011), Perfect Information Games with Upper Semicontinuous Payoffs, *Math. of OR*, **36(3)**, 468-473.
- [9] Shapley L.S. (1953), Stochastic Games, *Proceedings of National Academy of Science*, **39**, 1095-1100.
- [10] Solan E. and Vieille N. (2003), Deterministic Multi-Player Dynkin Games. *Journal of Mathematical Economics*, **39**, 911-929.
- [11] Thuijsman F. and Raghavan T.E.S. (1997), Perfect information stochastic games and related classes, *International Journal of Game Theory*, **26**, 403-408.
- [12] Ummels M. (2005), Rational Behavior and Strategy Construction in Infinite Multiplayer Games, Thesis, Rheinisch-Westf alischen Technischen Hochschule Aachen.
- [13] Vieille N. (2000a), Equilibrium in 2-person stochastic games I: A reduction. *Israel Journal of Mathematics*, **119**, 55-91.
- [14] Vieille N. (2000b), Equilibrium in 2-person stochastic games II: The case of recursive games. *Israel Journal of Mathematics*, **119**, 93-126.