



Deterministic n -person shortest path and terminal games on symmetric digraphs have Nash equilibria in pure stationary strategies

Endre Boros¹ · Paolo Giulio Franciosa² · Vladimir Gurvich^{1,3} · Michael Vyalyi^{3,4,5}

Accepted: 17 September 2023
© The Author(s) 2023

Abstract

We prove that a deterministic n -person shortest path game has a Nash equilibrium in pure and stationary strategies if it is edge-symmetric (that is (u, v) is a move whenever (v, u) is, apart from moves entering terminal vertices) and the length of every move is positive for each player. Both conditions are essential, though it remains an open problem whether there exists a NE-free 2-person non-edge-symmetric game with positive lengths. We provide examples for NE-free 2-person edge-symmetric games that are not positive. We also consider the special case of terminal games (shortest path games in which only terminal moves have nonzero length, possibly negative) and prove that edge-symmetric n -person terminal games always have Nash equilibria in pure and stationary strategies. Furthermore, we prove that an edge-symmetric 2-person terminal game has a uniform (subgame perfect) Nash equilibrium, provided any infinite play is worse than any of the terminals for both players.

Keywords Nash equilibrium · n -Person deterministic graphical games · Shortest path games · Terminal games

1 Introduction

Given a finite directed graph $G = (V, E)$, we interpret a vertex $v \in V$ as a position of a game and a directed edge $e = (u, v) \in E$ as a move from u to v .

The set of players is denoted by $I = \{1, \dots, n\}$. Each player $i \in I$ controls a subset $V_i \subseteq V$ of the positions, and we also have a nonempty subset V_T , the set of so-called *terminal* positions that are not controlled by any of the players. We assume that the sets V_1, \dots, V_n , and V_T form a partition of V , and that terminals are the only positions with no directed edges leaving. Furthermore, we fix a position $v_0 \in V \setminus V_T$ and call it the *initial* position.

Extended author information available on the last page of the article

A pure and stationary strategy σ_i of player $i \in I$ is a mapping $\sigma_i : V_i \mapsto V$ such that $(v, \sigma_i(v)) \in E$ for all $v \in V_i$. To simplify our notation, sometimes we also use σ_i as the set of edges $\{(v, \sigma_i(v)) \mid v \in V_i\}$. In this paper we consider only pure and stationary strategies, and call them simply strategies.

Let us denote by Σ_i the set of strategies of player $i \in I$, and set $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$. The tuple $\sigma = (\sigma_i \mid i \in I) \in \Sigma$ is called a *strategy profile* or *situation*. Every situation determines uniquely a directed walk, called the *play* $P(\sigma)$ and defined as follows. In a position $v \in V_i$ on this walk, player i makes the move $(v, \sigma_i(v))$, after which the walk enters $v' = \sigma_i(v)$. Starting with v_0 , this process creates a unique play $P(\sigma)$. This play either comes to a terminal $v \in V_T$, where it stops (no player handles the terminal and no moves are from it), or it repeats a position reached earlier, in which case it follows the enclosed cycle infinitely many times. In the first case we call $P(\sigma)$ a *terminal play*, while in the second case it is called an *infinite play*.

Each player $i \in I$ has his own length function $\ell^i : E \rightarrow \mathbb{R}$ that assigns to every move $e \in E$ a real number $\ell^i(e)$ that we interpret as the cost of this move for this player; we call it the ℓ^i -cost of the move and set $\ell = (\ell^i \mid i \in I)$. The cost of a sequence of moves is the sum of the costs of these moves. In particular, for a situation $\sigma \in \Sigma$ and player $i \in I$ the *effective cost* of σ for player i is

$$\ell^i(\sigma) = \sum_{(u,v) \in P(\sigma)} \ell^i(u, v).$$

This kind of additive effective cost function, called *total*, is considered in Thuijsman and Vrieze (1987); see more analysis and variations in Boros et al. (2018, 2017); Gurvich and Oudalov (2014). Note that the cost of an infinite play may not be well defined, in particular when the cost of moves can be both positive and negative. The above cited papers provide several possible solutions for this problem. In our paper we focus on the special case, when all cost functions are positive; in this case the cost of an infinite play is $+\infty$, while every terminal play has a finite cost.

All players aim to minimize their effective costs.

We will call the obtained class of games *the shortest path games*, or simply games in the sequel. Such a game Γ is defined by the triple (G, v_0, ℓ) . Note that we always assume a given partition $V = V_1 \cup \dots \cup V_n \cup V_T$ of the vertices of G ; just we do not want to convolute our paper with extra notation.

A situation $\sigma \in \Sigma$ is called a *Nash equilibrium* (or NE, in short) if $\ell^i(\sigma) \leq \ell^i(\sigma')$ for any player $i \in I$ and any situation $\sigma' \in \Sigma$ that may differ from σ only by the strategy of player i . In other words, a situation is a NE if no player $i \in I$ can reduce her effective cost by choosing another strategy, provided that all other players keep their old strategies. We call σ a *terminal NE* if it is a NE and $P(\sigma)$ is terminal.

1.1 Edge-symmetric and positive shortest path games

A local cost function ℓ is called *positive* if $\ell^i(e) > 0$ for all players $i \in I$ and moves $e \in E$. This condition is equivalent to the seemingly weaker condition of $\sum_{e \in C} \ell^i(e) > 0$ for all directed cycles C of G and players $i \in I$. The equivalence is based on the concept of potential transformation introduced for directed graphs in

Gallai (1958). It is applied to the length function of each player, separately. Such an equivalent transformation can in fact be computed efficiently, by linear programming. Since this transformation changes all path lengths between two vertices by the same constant, a shortest path game with positive cycle lengths can always be transformed into an equivalent game with positive edge lengths.

A digraph $G = (V, E)$ is called *edge-symmetric* if $(u, v) \in E$ whenever $(v, u) \in E$, except if $u \in V_T$. Finally, we call a game $\Gamma = (G, v_0, \ell)$ *edge-symmetric and positive* if G is edge-symmetric and ℓ is positive. Let us emphasize that the cost functions of the players are not assumed to be symmetric, that is $\ell^i(u, v) \neq \ell^i(v, u)$ may hold for some edges $(u, v) \in E$ and players $i \in I$.

Theorem 1 *An edge-symmetric and positive n -person shortest path game $\Gamma = (G, v_0, \ell)$ has a NE. Furthermore, it has a terminal NE whenever V_T is reachable from v_0 .*

Let us add that the above result cannot be extended to the so-called limit mean effective payoff games (see Gillette 1957 for definitions). For mean effective payoff a 2-person NE-free example was constructed in Gurvich (1988) on an edge-symmetric bipartite digraph with non-positive length function, and it is easy to make this example positive by adding a constant to the lengths of all edges.

Both conditions of edge-symmetry and positivity, are essential in Theorem 1. In Sect. 6 we provide examples of NE-free edge-symmetric (but not positive) 2-person games. For $n \geq 3$ such examples were known earlier (Boros and Gurvich 2003; Gurvich and Oudalov 2014). Furthermore, a NE-free positive but not edge-symmetric 3-person game was obtained in Gurvich and Oudalov (2014).

Another case, when the assumption of edge-symmetry helps to give criteria for the existence of a NE is the family of so-called cyclic games, in which each directed cycle is a separate outcome. For 2-person edge-symmetric cyclic games (Boros et al. 2011) provides a criterion for the existence of NE for all cost functions. Without the assumption of edge-symmetry no such criterion is known.

In Fraenkel et al. (1993), a polynomial algorithm solving game “geography” was found for symmetric digraphs, while the problem is known to be PSPACE-complete in general (Lichtenstein and Sipser 1980; Schaefer 1978).

Note also that positivity is important even if we have only one player. In the presence of negative arc lengths (with some negative cycles) a shortest path may not exist, and computing a shortest simple path is NP-hard in this case.

A shortest path game is called play-once if $|V_i| = 1$ for all $i \in I$. It was shown in Boros and Gurvich (2003) that every play-once positive game has a NE.

It remains an open problem whether every 2-person non-edge-symmetric positive shortest path game has a NE or not.

1.2 Edge-symmetric terminal games

A related family of games that in fact forms a special case of shortest path games is the family of terminal games. In many popular positional games (Go,

Chess, Checkers, etc.) players do not pay for moves; the effective cost for each player depends only on the terminal position. We assume, as before, that the set of positions and moves of the game form a directed graph, $G = (V, E)$, players $i \in I$ control distinct subsets $V_i \subseteq V$ of the positions, and V_T is a nonempty subset of the terminal positions, such that $V = V_1 \cup \dots \cup V_n \cup V_T$ is a partition of the set of positions. In these games a situation may end up in a terminal position within V_T or in an infinite play, going around the same cycle infinitely many times. We consider all infinite plays equivalent, and denote the corresponding outcome by c . Thus, to evaluate the outcome of these games we have mappings $\mathcal{L}^i : V_T \rightarrow \mathbb{R}$ for $i \in I$ that evaluate the value of terminals for players.

Then, for a situation σ that defines a play $P(\sigma)$ the effective cost of player $i \in I$ is defined as

$$\mathcal{L}^i(\sigma) = \begin{cases} \mathcal{L}^i(w) & \text{if } P(\sigma) \text{ is a finite play terminating at } w \in V_T, \\ 0 & \text{if } P(\sigma) \text{ is the infinite play } c. \end{cases}$$

The infinite play c may be worse than some terminals and better than some others for a player. As in shortest path games, all players minimize their own effective costs. We call such a game $\Gamma = (G, v_0, \mathcal{L})$ a *terminal game*.

In the 2-person zero-sum case, such games are also called *deterministic graphical*; see (Washburn 1990), where the existence of a NE (saddle point) in pure and stationary strategies was shown.

In Boros and Gurvich (2003), this result was extended to the non-zero-sum 2-person case using a general criterion of Gurvich (1975, 1988). Recently these results were extended further for the so-called multi-stage deterministic graphical games; see (Gurvich 2018) and also (Gurvich and Koshevoy 2018). However, this line of arguments cannot be extended to the n -person case for $n > 2$, since the criterion of Gurvich (1975, 1988) holds only for $n = 2$. Examples for 3- and 4-person NE-free terminal games were constructed in Boros et al. (2018); Gurvich (2015).

Let us also mention that Everett (1997) already in 1957 considered a closely related class of terminal games and, among other results, provided some NE-free examples for concurrent zero-sum terminal games.

The existence of a NE for the above defined class of terminal games remains an open problem if we assume additionally that $\mathcal{L}^i(v) < 0$ for all terminals $v \in V_T$ and players $i \in I$. This condition describes a natural subclass of these games in which an infinite play is worse than any terminal play for each of the players. We can show that this subfamily of terminal games can be viewed as a subfamily of positive shortest path games, and thus the existence of a (terminal) NE for such a game is implied by Theorem 1. We can further generalize this by waiving the above assumption, and prove the following claim:

Theorem 2 *An edge-symmetric n -person terminal game has a NE.*

1.3 Uniform Nash equilibrium for edge-symmetric terminal games

We can further strengthen Theorem 2 for the case of 2 players. We call a situation $(\sigma_i \mid i \in I)$ a *uniform Nash equilibrium* (or a UNE, in short), if it is a NE no matter what the initial position is.

Theorem 3 *Assume that $\Gamma = (G, \mathcal{L})$ is a terminal game that satisfies the following conditions:*

- (TWO) $n = |I| = 2$;
- (SYM) G is edge-symmetric;
- (CIW) *infinite plays are worse than any of the terminal plays for all players.*

Then Γ has a UNE.

Let us add that all three conditions are necessary to guarantee the existence of a UNE, as we demonstrate it in Sect. 5.

We can also note that condition (CIW) is automatically satisfied by positive shortest path games. However, the theorem still fails even if we keep all three conditions but replace terminal games by shortest path games, see the last example in Sect. 6.

2 Edge-symmetric positive shortest path games

Assume that $\Gamma = (G, v_0, \ell)$ is an edge-symmetric and positive shortest path game. We denote by $N^+(v) = \{u \mid (v, u) \in E\}$ the out-neighborhood of a position $v \in V$ (and thus we have $N^+(v) = \emptyset$ for all terminals $v \in V_T$). Note also that since all plays reaching a terminal end there, we can merge all terminals into a single terminal position without any loss of generality. Thus, in this section we assume that $V_T = \{v_t\}$.

Let us consider the subgraphs induced by vertex sets $V_i, i \in I$. Each such subgraph can uniquely be decomposed into strongly connected subgraphs. We denote by $\mathcal{Q} = \{Q_1, Q_2, \dots\}$ the family of all such components for all $i \in I$. For sake of simplicity, we consider $\{v_t\}$ also as one of those components. Note that we adopt the name “component” for these subgraphs, even though they may not be strongly connected components of our graph. They are strongly connected components of the subgraphs, induced by $V_i, i \in I$ and $V_T = \{v_t\}$.

We use $i(v)$ to denote the player who controls vertex v . We extend this notation even for the terminal node for notational convenience and assume $i(v_t) \notin I$, even though v_t has no outgoing arcs, and no player controls it. We denote by $i(Q) \in I$ the player who controls the vertices in such a component $Q \in \mathcal{Q}$. Note that, due to the symmetric nature of the graph and the above definition of the components, we have the following property:

- (A) If $u \in Q, v \in Q', Q \neq Q'$ and $(u, v) \in E$ then $i(Q) \neq i(Q')$.

Since a component $Q \in \mathcal{Q}$ is strongly connected, for any two vertices $u, v \in Q$ player $i(Q)$ can create directed path(s) between these vertices. We denote by $d^{i(Q)}(u, v)$ the length of a shortest $u \rightarrow v$ path within Q , where the $i(Q)$ -length is used to measure the length of the edges. Note that we do not assume the symmetry of the lengths, thus $d^{i(Q)}(u, v)$ and $d^{i(Q)}(v, u)$ may be different values.

Let us define $Q(v) \in \mathcal{Q}$ to be the component containing vertex $v \in V$.

For a $v_0 \rightarrow v_t$ path P let us call a sequence $\{v_i, v_{i+1}, \dots, v_{i+j}\}$ a subpath of P . The sets $P \cap Q$, $Q \in \mathcal{Q}$ partition P into a number of such subpaths. We denote by $q(P)$ the number of these subpaths. In other words, if we denote by $Q(P)_j$, $j = 1, \dots, q$ the components that P intersects as we follow P from v_0 to v_t , then $q = q(P)$, and $Q(P)_j \neq Q(P)_{j+1}$ for all $j = 1, \dots, q - 1$. Note that we may have $Q(P)_j = Q(P)_k$ for some j and $k \geq j + 2$.

2.1 Special paths

Let us now consider a $v_0 \rightarrow v_t$ path P in G , a player $i \in I$, and the subgraph $G(i, P)$ of G consisting of the moves in P and the moves $(u, v) \in E$ with $u \in V_i$. We say that P is *i-special* if P is a shortest $v_0 \rightarrow v_t$ path in $G(i, P)$ for player i .

Lemma 1 *There exists a $v_0 \rightarrow v_t$ path P satisfying the following conditions:*

- (B) $q(P)$ is the smallest among all $v_0 \rightarrow v_t$ paths, and
- (C) P is *i-special* for all $i \in I$.

Proof Let us first associate a new edge length $\lambda : E \mapsto \{0, 1\}$ by defining

$$\lambda(u, v) = \begin{cases} 1 & \text{if } Q(u) \neq Q(v), \\ 0 & \text{if } Q(u) = Q(v), \end{cases}$$

and choose a λ -shortest $v_0 \rightarrow v_t$ path P . By abusing our notations, we shall use P to denote both the set of vertices of P and also the set of edges in P . We hope that from the context it will always be unambiguous what we mean.

Note that $q(P)$ is minimum among all $v_0 \rightarrow v_t$ paths, and that P does not cross any component $Q \in \mathcal{Q}$ twice. This is because Q is strongly connected and the λ -cost of a path inside Q is zero, and outside is positive. Let us observe that the minimality of $q(P)$ implies the following property:

- (D) For all $v \in Q(P)_j$, $1 \leq j \leq q(P) - 2$ and $j + 2 \leq k \leq q(P)$ we have $N^+(v) \cap Q(P)_k = \emptyset$.

Let us introduce $P[j] = P \cap Q(P)_j$ for $j = 1, \dots, q(P)$ and note that the edges in P leaving vertices in $P[j]$ form a subpath ending in $Q(P)_{j+1}$. Let us next define

$$r(P)_j = \sum_{\substack{u \in P[j] \\ (u, v) \in P}} e^{i(Q(P)_j)}(u, v)$$

for $j = 1, \dots, q(P) - 1$, and set $r(P) = (r(P)_1, \dots, r(P)_{q(P)-1})$. Note that in a way $r(P)$ measures the length of P such that the length of every move is measured by the length function of the player who controls that move.

Assume now that P is an arbitrary path that satisfies property (B). Our arguments above show that there are such paths. Assume also that it does not satisfy property (C), that is there exists a player $i \in I$ that can improve on it. It means that there is an index $1 \leq j < q(P)$ with $i = i(Q(P)_j)$ and a position $v \in P[j]$ such that player i can deviate from P at v , return to P in $Q(P)_j \cup Q(P)_{j+1}$ (due to property (D)) and make the i -length of the improved path P' shorter than the i -length of P . Note that if P' returns to P at or before the first position of $P[j + 1]$, then we must have $r(P')_j < r(P)_j$ and $r(P')_k = r(P)_k$ for all $k > j$. If P' returns to P at or after the second position of $P[j + 1]$ (in case $P[j + 1]$ has more than one vertex), then we must have $r(P')_{j+1} < r(P)_{j+1}$ and $r(P')_k = r(P)_k$ for all $k > j + 1$. Note that in the second case we may have $r(P')_j > r(P)_j$. In both cases however, the vector $r(P')$ is smaller than $r(P)$ in the reverse lexicographic order. Note furthermore that P' must also satisfy property (B).

Since we have only finitely many different paths, and thus $r(P)$ vectors, after finitely many improvement steps we must arrive to a path that satisfies both (B) and (C), as claimed. \square

Let us call a path P satisfying both (B) and (C) a *special path*. By the above lemma such a special path exists. Let us now fix one $v_0 \rightarrow v_1$ special path P .

2.2 Extending a special path to a NE

Let us introduce for $j = 1, \dots, q(P)$ the sets

$$U(P)_j = \bigcup_{k=1}^j Q(P)_k,$$

and let $u_j \in Q(P)_j$ be the first vertex on $P \cap Q(P)_j$. Recall that for a vertex $u \in Q(P)_j$ and player $i = i(Q(P)_j)$ we denote by $d^i(u_j, u)$ the i -distance from u_j to u inside component $Q(P)_j$.

We are ready now to define strategies for the players.

- (i) If $v \in P$ then we choose edge $(v, u) \in P$.
- (ii) If $v \in V \setminus P$ and $N^+(v) \cap U(P)_{q(P)} \neq \emptyset$, then
 - (ii-1) choose the smallest index k such that $N^+(v) \cap Q(P)_k \neq \emptyset$ and set $i = i(Q(P)_k)$;

- (ii-2) choose a vertex $u \in N^+(v) \cap Q(P)_k$ that minimizes $d^i(u_k, u)$;
(ii-3) choose edge (v, u) .
- (iii) If $v \in V \setminus P$ and $N^+(v) \cap U(P)_{q(P)} = \emptyset$, then choose an arbitrary edge $(v, u) \in E$

Let us denote by $\sigma(P) = \{\sigma_i(P) \mid i \in I\}$ one of the situations defined in this way.

Lemma 2 *If P is a special $v_0 \rightarrow v_t$ path, then $\sigma(P)$ is a NE.*

Proof Assume for contradiction that a player $i \in I$ can deviate and improve on P . This means that there is a component $Q(P)_j$ controlled by $i = i(Q(P)_j)$ and there is a vertex in $P \cap Q(P)_j$ such that player i can deviate from P starting at this vertex and create a new path P' such that the i -length of P' is smaller than the i -length of P . We can assume that P' is the shortest path (according to ℓ^i) player i can create, and that $P \cap Q(P)_k = P' \cap Q(P)_k$ for all $k < j$.

In particular, P' cannot return to $P \cap U(P)_{j-1}$. Note first that if $u \in U(P)_{j-1} \setminus P$, $i(u) \neq i$, and $(u, v) \in \sigma(P)$, then we must have $v \in U(P)_{j-1}$ by our definition of σ , since our graph is edge-symmetric. Note next that if $u \in U(P)_{j-1}$, and $i(u) = i$, then $u \in U(P)_{j-2}$. These together imply that if P' enters $U(P)_{j-1}$ then it cannot reach v_t .

Let us now focus on the segment of P' that is outside of $U(P)_{j-1}$ (note that this may be the entire path P' if $j = 1$.) Let us denote the vertices along this segment of P' by $u_j = w_0, w_1, \dots, w_m = v_t$. Let k be the smallest index such that $w_k \notin Q(P)_j$. By property (C) and the fact that P' is the shortest path that player i can create, we can conclude that $w_k \notin P$. Then we must have $i(w_k) \neq i$ by the definition of the components, thus we have $(w_k, w_{k+1}) \in \sigma(P)$. By the symmetric nature of our graph, by part (ii) of our definition of σ , and by the fact that P' does not enter $U(P)_{j-1}$ we must have $w_{k+1} \in Q(P)_j$. Since $N^+(w_k) \cap Q(P)_j \supseteq \{w_{k-1}, w_{k+1}\}$, we get that (ii-2) of the definition of σ implies $d^i(u_j, w_{k+1}) \leq d^i(u_j, w_{k-1})$. Since player i can create such a shortest $u_j \rightarrow w_{k+1}$ path inside $Q(P)_j$ he could replace w_0, \dots, w_{k+1} with this path, and create another path P'' that is strictly shorter (by at least $\ell^i(w_{k-1}, w_k) + \ell^i(w_k, w_{k+1})$) than P' , contradicting the fact that P' is the shortest path this player can create deviating from P inside $Q(P)_j$. This contradiction proves our claim. \square

Proof of Theorem 1 The claim follows by Lemmas 1 and 2. \square

3 Terminal games and shortest path games

It is easy to see that terminal games satisfying condition (CIW) of Theorem 3 can also be viewed as positive shortest path games.

Theorem 4 *Assume that $\Gamma = (G, v_0, \mathcal{L})$ is a terminal game that satisfies condition (CIW). Then there exist positive local costs functions $\ell^i : E \mapsto \mathbb{R}$ for $i \in I$ such that*

$\Gamma' = (G, v_0, \ell)$ is a positive shortest path game such that any situation σ that is a terminal NE of Γ' is also a NE of Γ .

Proof Let us note first that condition (CIW) implies that $\mathcal{L}^i(v) < 0$ for all terminals $v \in V_T$ and players $i \in I$. Since we have only finitely many terminals and players, we can assume w.l.o.g. that $\mathcal{L}^i(v)$ are all negative integers for all terminals $v \in V_T$ and players $i \in I$, and there exists a positive integer M such that

$$-M < \mathcal{L}^i(v) < 0 \quad \text{for all } i \in I, \text{ and } v \in V_T.$$

Let us next define the local costs for the players $i \in I$ and moves $(u, v) \in E$ by

$$\ell^i(u, v) = \begin{cases} \frac{1}{2|E|} & \text{for all moves } (u, v) \in E, v \notin V_T, \\ M + \mathcal{L}^i(v) & \text{for all moves } (u, v) \in E, v \in V_T. \end{cases}$$

Thus, for an arbitrary situation σ for which $P(\sigma)$ is a finite play terminating at $v \in V_T$ and for a player $i \in I$ we have

$$\ell^i(\sigma) = \sum_{(u,v) \in P(\sigma)} \ell^i(u, v) = \frac{|P(\sigma)| - 1}{2|E|} + (M + \mathcal{L}^i(v))$$

implying that for all situations σ with a finite play and for all players $i \in I$ we have

$$M + \mathcal{L}^i(\sigma) \leq \ell^i(\sigma) < M + \mathcal{L}^i(\sigma) + \frac{1}{2}.$$

Now assume that σ is a terminal NE in Γ' and let σ' be obtained from σ by one of the players, say $i \in I$ changing his strategy. If σ' is infinite, then $\mathcal{L}^i(\sigma') > \mathcal{L}^i(\sigma)$. Otherwise, by the above inequalities we can write

$$M + \mathcal{L}^i(\sigma) \leq \ell^i(\sigma) \leq \ell^i(\sigma') < M + \mathcal{L}^i(\sigma') + \frac{1}{2},$$

where the second inequality follows from the fact that σ is NE in Γ' . Since the \mathcal{L}^i values are integers, $\mathcal{L}^i(\sigma) \leq \mathcal{L}^i(\sigma')$ follows, proving that σ is also a NE in Γ . \square

Note, that in the above construction, Γ may have some NE that are not NE in Γ' .

Note also that Γ and Γ' in the above statement use the same underlying directed graph. Thus they are simultaneously edge-symmetric or non-edge-symmetric and we can derive the following claim:

Corollary 5 *Edge-symmetric terminal games satisfying condition (CIW) have NE. Furthermore, there is a terminal NE whenever V_T is reachable from v_0 .*

Proof By Theorem 4, an edge-symmetric terminal game satisfying condition (CIW) can be viewed as a positive shortest path game, and thus the claim follows by Theorem 1. \square

4 Edge-symmetric terminal games

Let us consider an edge-symmetric terminal game $\Gamma = (G, v_0, \mathcal{L})$, and note that unlike in shortest path games, a loop (that is a move of the form (u, u) for some position $u \in V$) may play a role in a NE, since the infinite play may not be the worse outcome for some players. Consequently, we have $u \in N^+(u)$ whenever $(u, u) \in E$.

Analogously to shortest path games, let us consider the family \mathcal{Q} of strongly connected components of the subgraphs induced by the subsets $V_i, i \in I$.

We associate to Γ its *small* version Γ' obtained by “merging” the strongly connected components $Q \in \mathcal{Q}$ into single positions. More precisely, we introduce a new position v_Q and set $i(v_Q) = i(Q)$ for all $Q \in \mathcal{Q}$, and define a directed graph G' on the set of positions $V' = V_T \cup \{v_Q \mid Q \in \mathcal{Q}\}$. We define the edge set E' of G' by including edges (v_Q, v_R) for $Q, R \in \mathcal{Q}, Q \neq R$ if there are positions $u \in Q$ and $v \in R$ such that $(u, v) \in E$. We also include edge (v_Q, w) for $Q \in \mathcal{Q}$ and $w \in V_T$ if there is a position $u \in Q$ such that $(u, w) \in E$. Furthermore, we include a looping edge (v_Q, v_Q) for all $Q \in \mathcal{Q}$ with $|Q| \geq 2$. Furthermore, if $Q = \{u\}$ for some $Q \in \mathcal{Q}$ and $(u, u) \in E$, then we also include in E' the loop (v_Q, v_Q) . Note that with the above definitions we avoided creating parallel edges, since they would be redundant in a terminal game. Finally, we define $v'_0 = v_Q$ if $v_0 \in Q$. Note that this small game $\Gamma' = (G', v'_0, \mathcal{L})$ is again an edge-symmetric terminal game with the same set of players and cost function as Γ .

Lemma 3 *If the small game Γ' has a NE then so does Γ .*

Proof Assume that $\sigma' \subseteq E'$ is a NE of Γ' . We associate to σ' a situation σ in Γ as follows.

For every position $v_Q \in V', Q \in \mathcal{Q}$ there exists a unique edge $(v_Q, x) \in \sigma'$. We associate to (v_Q, x) edges of E to be included in σ in the following way:

If $x = v_R \neq v_Q$, then by the definition of the small game we have (at least one) corresponding edge $(u, v) \in E$ such that $u \in Q, v \in R$. Let us consider one of these (u, v) edges, and declare position $u = u^Q$ the root of the strong component $Q \in \mathcal{Q}$. Since $Q \in \mathcal{Q}$ is a strong component, we have a directed tree T^Q rooted at $u^Q \in Q$ such that from every vertex $u' \in Q$ we have a unique directed path from u' to u^Q via the edges of T^Q . Let us then include in σ the edges of T^Q and edge $(u, v) = (u^Q, v)$.

If $x = v_Q$ and $|Q| \geq 2$, then (by the definition of a component) Q is a strongly connected subgraph with at least two vertices, $u, v \in Q, u \neq v$. Let us declare $u = u^Q$ the root of Q and consider a directed tree T^Q on the vertices of Q such that from all other vertices of Q there is a unique path to u^Q . Let us then include in σ the edges of T^Q and the edge (u^Q, v) .

Finally, if $x = v_Q$ and $Q = \{u\}$, then by the definition of the small game we must have $(u, u) \in E$, and we include this loop in σ .

Note that we have $\mathcal{L}^i(\sigma) = \mathcal{L}^i(\sigma')$ for all $i \in I$. Furthermore, for all positions $v \in Q \in \mathcal{Q}$ player $i(v) = i(Q)$ can reach the same set of terminals (and/or the infinite play) in Γ and Γ' , assuming all other players keep their strategies. Thus, since σ' is a NE in Γ' , situation σ must also be a NE in Γ . \square

Proof of Theorem 2 Assume that $\Gamma = (G, v_0, \mathcal{L})$ is an edge-symmetric terminal game. By Lemma 3 we can assume that $|Q| = 1$ for all $Q \in \mathcal{Q}$, or in other words that for all moves $(u, v) \in E$ we have $i(u) \neq i(v)$, unless $u = v$.

Clearly, if V_T is not reachable from v_0 then any situation is a NE. Assume for the rest of our proof that V_T is reachable from v_0 .

For all positions $v \in V$ with $N^+(v) \cap V_T \neq \emptyset$ let us denote by $t(v) \in V_T$ a terminal that is a most preferred by player $i(v)$ in $N^+(v) \cap V_T$.

We consider the following cases:

Case 1: $N^+(v_0) \cap V_T = \emptyset$ and $\exists v_1 \in N^+(v_0)$ such that $N^+(v_1) \cap V_T = \emptyset$.

In this case we can construct a NE σ resulting in an infinite play. We include in σ the moves (x, v_0) for all $x \in N^+(v_0)$, (v_0, v_1) , and (y, v_1) for all $y \in N^+(v_1) \setminus (N^+(v_0) \cup \{v_0\})$. For all other positions we include an arbitrary move. Here $P(\sigma)$ is the infinite play $v_0 \rightarrow v_1 \rightarrow v_0$, and neither $i(v_0)$ nor $i(v_1)$ can achieve a better outcome. Note that $v_0 = v_1$ is possible in this case.

Case 2: $N^+(v_0) \cap V_T = \emptyset$ and $N^+(v) \cap V_T \neq \emptyset$ for all $v \in N^+(v_0)$.

Let us define $v_1 \in N^+(v_0)$ as a position such that $t(v_1)$ is one of the most preferred terminals for $i(v_0)$ in the subset $\{t(v) \mid v \in N^+(v_0)\} \subseteq V_T$.

Subcase 2.1: Player $i(v_1)$ prefers $t(v_1)$ to an infinite play.

We can construct a NE σ in this case as follows. We include the move (v_0, v_1) , the moves (x, v_1) for all $x \in N^+(v_1) \setminus V_T$, the moves $(v, t(v))$ for all $v \in N^+(v_0) \setminus N^+(v_1)$, and arbitrary moves for all other positions. Now the play $P(\sigma)$ is the path $v_0 \rightarrow v_1 \rightarrow t(v_1)$ yielding $\mathcal{L}^i(\sigma) = \mathcal{L}^i(t(v_1))$ for all players $i \in I$. Only two players are involved in the play, $i(v_0)$ and $i(v_1)$. If $i(v_1)$ deviates from σ then he can get either another terminal in $N^+(v_1) \cap V_T$ or the infinite play, and neither one is better for him than $t(v_1)$. If $i(v_0)$ deviates then he can get one of the terminals $\{t(v) \mid v \in N^+(v_0)\}$, and by our choice, none of them is better for him than $t(v_1)$.

Subcase 2.2: For player i , the cost of $t(v_1)$ is at least that of the infinite play.

We can construct a NE σ in this case as follows. We include the moves (x, v_0) for all $x \in N^+(v_0)$, the moves (y, v_1) for all $y \in N^+(v_1) \setminus V_T$, and an arbitrary move for all other positions. Now $P(\sigma)$ is the infinite play $v_0 \rightarrow v_1 \rightarrow v_0$. Player $i(v_0)$ can only deviate to another infinite play. Player $i(v_1)$ can deviate to either another infinite play, or to a terminal in $N^+(v_1) \cap V_T$, but those are not better for him than $t(v_1)$, which is not better in this case than an infinite play.

Case 3: $N^+(v_0) \cap V_T \neq \emptyset$ and $N^+(v_0) \not\subseteq V_T$.

Let us now create a smaller subgame $\Gamma' = (G', v_0, \mathcal{L})$ by deleting from Γ the moves (v_0, w) for $w \in N^+(v_0) \cap V_T$. Since Case 1 or 2 holds for Γ' , we have a NE σ' in Γ' by the previous arguments.

Subcase 3.1: $\mathcal{L}^{i(v_0)}(\sigma') \leq \mathcal{L}^{i(v_0)}(t(v_0))$.

In this case σ' is also a NE in Γ , since only player $i(v_0)$ gained extra options, and all of those options would take him to a terminal in $N^+(v_0) \cap V_T$, which are assumed to be not better in this case for $i(v_0)$ than the outcome of σ' .

Subcase 3.2: $\mathcal{L}^{i(v_0)}(\sigma') > \mathcal{L}^{i(v_0)}(t(v_0))$.

Let us modify σ' by replacing the move from v_0 by $(v_0, t(v_0))$, and denote by σ the obtained situation of Γ . We claim that σ is a NE in Γ in this case. This is because

$P(\sigma)$ is the single move $v_0 \rightarrow t(v_0)$, and any deviation from σ by player $i(v_0)$ yields either another terminal in $N^+(v_0) \cap V_T$, which is not better than $t(v_0)$, or it yields σ' or a situation obtainable by a deviation from σ' in Γ' . Since σ' is a NE in Γ' none of these can give a better outcome to player $i(v_0)$ than σ' which is not better than $t(v_0)$ by our assumptions in this case.

Case 4: $N^+(v_0) \subseteq V_T$.

In this case choosing $\sigma_{i(v_0)}(v_0) = t(v_0)$ trivially gives a NE.

□

5 Existence of UNE in edge-symmetric terminal games

In this section we prove Theorem 3 claiming that under conditions (TWO), (SYM), and (CIW) a terminal game has a UNE.

We consider an edge-symmetric terminal game $\Gamma = (G, \mathcal{L})$ in which no initial position is fixed, and assume that $I = \{1, 2\}$. For a situation σ and arbitrary position $v \in V$ we denote by $P(\sigma, v)$ the unique walk starting from position v and following the moves in situation σ . Such a walk either terminates in a terminal position $w = w(\sigma, v) \in V_T$, or is infinite, cycling around a directed cycle infinitely many times. Accordingly, we define the effective cost for player $i \in I$ for a position $v \in V$ by

$$\mathcal{L}^i(\sigma, v) = \begin{cases} \mathcal{L}^i(w(\sigma, v)) & \text{if } P(\sigma, v) \text{ is terminating, and} \\ 0 & \text{if } P(\sigma, v) \text{ is infinite.} \end{cases}$$

Note that the above definition and condition (CIW) imply that we have $\mathcal{L}^i(w) < 0$ for all players $i \in I$ and terminals $w \in V_T$.

Given a player $i \in I$ we use σ^{-i} to denote the strategy of the opponent. Thus, we can write $\sigma = (\sigma^{-i}, \sigma_i)$ for an arbitrary situation and player $i \in I$, where $\sigma_i \in \Sigma_i$ is a strategy for player $i \in I$. Note that in this section we have $|I| = 2$, and thus for a situation $\sigma = (\sigma_1, \sigma_2)$ we have $\sigma^{-1} = \sigma_2$ and $\sigma^{-2} = \sigma_1$.

Given a player $i \in I$ and σ^{-i} a strategy $\sigma_i \in \Sigma_i$ is called a *uniform best response* of player i to σ^{-i} if the equality

$$\mathcal{L}^i((\sigma^{-i}, \sigma_i), v) = \min_{\sigma'_i \in \Sigma_i} \mathcal{L}^i((\sigma^{-i}, \sigma'_i), v)$$

holds for all positions $v \in V$. Note that a uniform best response always exists, can be determined efficiently, and may not be unique.

We call a situation $\sigma = (\sigma_i \mid i \in I)$ a *uniform Nash equilibrium* (or UNE, in short), if σ_i is a uniform best response of player i to σ^{-i} for all players $i \in I$.

Given a situation $\sigma = (\sigma^{-i}, \sigma_i)$ we say that strategy σ'_i is a *uniform best improvement* for player i if σ'_i is a uniform best response to σ^{-i} , different from σ_i , and for all positions $v \in V_i$ we have either $\sigma_i(v) = \sigma'_i(v)$ or $\mathcal{L}^i((\sigma^{-i}, \sigma'_i), v) < \mathcal{L}^i((\sigma^{-i}, \sigma_i), v)$. Note that every uniform best improvement is a uniform best response, but not necessarily the other way around. Furthermore, if a situation is not a UNE, then at least

one of the players have a uniform best improvement (and it can be determined efficiently). Let us also note that a uniform best improvement may not be unique in case players have ties over the set of terminals.

Our first step is to show that it is enough to prove Theorem 3 for simplified, special edge-symmetric terminal games.

Lemma 4 *If $\Gamma = (G, \mathcal{L})$ is an edge-symmetric terminal game, then we can assume w.l.o.g. that edges are between different vertices controlled by different players, i.e.,*

$$(u, v) \in E, v \notin V_T \implies i(u) \neq i(v). \tag{1}$$

We can also assume that from each position we have at most one terminal move, i.e.,

$$|N^+(v) \cap V_T| \leq 1 \quad \text{for all } v \in V. \tag{2}$$

We can assume further that the terminals are reachable from each position $v \in V$, i.e.,

$$\text{for all } v \in V \text{ there exist a } v \rightarrow V_T \text{ path.} \tag{3}$$

Proof For condition (1) let us also note that Lemma 3 with the same proof works also for UNE (instead of NE).

For the inequalities (2) note first that if $x, y \in N^+(v) \cap V_T, x \neq y$, and $\mathcal{L}^{i(v)}(x) < \mathcal{L}^{i(v)}(y)$ then the move (v, y) can not belong to a UNE. Furthermore, if $\mathcal{L}^{i(v)}(x) = \mathcal{L}^{i(v)}(y)$ and σ' is a UNE in Γ' obtained from Γ by deleting move (v, y) , then σ' is also a UNE in Γ . This is because only player $i(v)$ has an option in Γ that is not available in Γ' (namely the move (v, y)) but that move cannot be part of a uniform best improvement of player $i(v)$ to $\sigma'^{-i(v)}$ since the move (v, x) is available for player $i(v)$ in Γ' , σ' is a UNE in Γ' , and $\mathcal{L}^{i(v)}(x) = \mathcal{L}^{i(v)}(y)$.

For property (3) let us consider the set of positions $U \subseteq V$ such that no directed path connects a vertex $u \in U$ to the set of terminals V_T . Then, assigning arbitrary moves to positions in U does not change whether a situation is a UNE or not. Since identifying the set of positions from which the set of terminals is not reachable is a computationally easy task, and since those positions have no influence on the existence of a UNE, we can assume w.l.o.g. that we preprocess the input game, and delete all such positions, before any further analysis. \square

Lemma 5 *Assume that $\Gamma = (G, \mathcal{L})$ is an edge-symmetric terminal game, satisfying conditions (1) and σ is a situation such that for one of the players, $i \in I$, strategy σ_i is a uniform best response to σ^{-i} . Assume further that $u \in V_j, v \in V_i$ are distinct positions with $j \neq i$ that are connected by edges in G . Then we have*

$$\mathcal{L}^i(\sigma, v) \leq \mathcal{L}^i(\sigma, u).$$

Proof Let us note first that if (v, u) belongs to σ_i or $P(\sigma, u)$ contains position v , then we must have $\mathcal{L}^i(\sigma, v) = \mathcal{L}^i(\sigma, u)$ by the definition of a terminal game.

Assume next that the move (v, u) does not belong to σ_i and position v does not belong to $P(\sigma, u)$. Let us define a new strategy σ'_i for player i by

$$\sigma'_i(w) = \begin{cases} \sigma_i(w) & \text{for all } w \in V_i \setminus \{v\}, \\ u & \text{for } w = v. \end{cases}$$

With this definition we have

$$\mathcal{L}^i(\sigma, u) = \mathcal{L}^i((\sigma^{-i}, \sigma'_i), u) = \mathcal{L}^i((\sigma^{-i}, \sigma'_i), v) \geq \mathcal{L}^i((\sigma^{-i}, \sigma_i), v).$$

Here the first two equalities follow by the definition of σ'_i and the fact that position v does not belong to $P(\sigma, u)$. The last inequality follows by the assumption that σ_i is a uniform best response to σ^{-i} . \square

Lemma 6 *Assume that $\Gamma = (G, \mathcal{L})$ is a terminal game satisfying condition (CIW), and σ is a situation such that the plays $P(\sigma, v)$ are finite for all $v \in V$. Then, if σ' is obtained from σ by a uniform best improvement of one of the players, the plays $P(\sigma', v)$, $v \in V$ are also all finite.*

Proof This is an immediate consequence of condition (CIW) and the definition of uniform best improvement. \square

Our proof of Theorem 3 is based on an iterative process, the finiteness of which depends on a strictly monotone decreasing measure of progress. The proof of that monotonicity depends critically on the following definition.

Let us assume, for the rest of this section, that by (2) of Lemma 4 we have $|N^+(v) \cap V_T| \leq 1$ for all $v \in V$. Let us then denote by $t(v) \in N^+(v) \cap V_T$ the unique terminal for vertices $v \in V$ with $N^+(v) \cap V_T \neq \emptyset$.

Let us call a situation σ *i-basic* for a player $i \in I$ if for all positions $v \in V$ with $N^+(v) \cap V_T \neq \emptyset$ we either have $\sigma(v) = t(v)$ or $\mathcal{L}^{i(v)}(\sigma', v) < \mathcal{L}^{i(v)}(t(v))$ hold for all situations σ' obtained from σ by a uniform best improvement of player i . In other words, a situation is *i-basic* if a uniform best improvement by player i provides every position that it controls and has an unused terminal move with a strictly better outcome than that terminal move would provide.

Lemma 7 *Assume that $\Gamma = (G, \mathcal{L})$ is an edge-symmetric terminal game that satisfies the conditions in Lemma 4. Then we can efficiently find a situation σ that is *i-basic* for all players $i \in I$, and for which the plays $P(\sigma, v)$ are finite for all $v \in V$.*

Proof We can construct a situation σ satisfying the claimed properties by the following approach. Let us define $W = \{v \in V \mid N^+(v) \cap V_T \neq \emptyset\}$ and note that by our assumptions we have a terminal $t(v) \in V_T$ for all $v \in W$ such that $N^+(v) \cap V_T = \{t(v)\}$ holds. Let us now choose the moves $(v, t(v))$ for all positions $v \in W$, and color all positions in $W \cup V_T$ blue. In the sequel, while there exists an uncolored position $v \in V$ from which we can reach a blue position u in one move, we choose the move (v, u) and color v blue. In a finite number of steps this procedure will stop, and all positions will be colored blue, since we assume that there is a finite directed path from all positions to the set of terminals. In this way we define a situation σ such that all plays $P(\sigma, v)$ are finite.

To see that σ is i -basic, for all players $i \in I$, let us consider one of the players i , and derive σ' by applying a uniform best improvement of player i to σ . Now, for a position $v \in W \setminus V_i$ we have $\sigma'(v) = \sigma(v) = t(v)$ by our definitions of σ and σ' , while for a position $v \in W \cap V_i$ we have either $\sigma'(v) = \sigma(v) = t(v)$ or $\mathcal{L}^i(\sigma', v) < \mathcal{L}^i(\sigma, v) = \mathcal{L}^i(t(v))$, by the definition of a uniform best improvement. \square

Lemma 8 *Assume that $\Gamma = (G, \mathcal{L})$ is a 2-person edge-symmetric terminal game satisfying conditions (CIW), conditions in Lemma 4, and that σ is a 1-basic situation such that all plays $P(\sigma, v)$ are finite. Assume further that σ^1 is obtained from σ by a uniform best improvement of player 1, and σ^2 is obtained from σ^1 by a uniform best improvement of player 2. Then we have the inequalities*

$$\mathcal{L}^{i(v)}(\sigma^2, v) \leq \mathcal{L}^{i(v)}(\sigma^1, v)$$

for all positions $v \in V$.

Proof Since σ^2 is obtained from σ^1 by a uniform best improvement of player 2, the claim follows for all $v \in V_2$ by the definition of a uniform best improvement.

Let us next consider positions $v \in V_1$. If $P(\sigma^1, v) = P(\sigma^2, v)$, then we have the equality $\mathcal{L}^1(\sigma^2, v) = \mathcal{L}^1(\sigma^1, v)$.

Assume for the rest of our proof that $P(\sigma^2, v) \neq P(\sigma^1, v)$. Since $P(\sigma^2, v)$ is a finite play by Lemma 6, positions controlled by player 1 and 2 alternate along this finite path, due to our assumption (1). Since $P(\sigma^2, v) \neq P(\sigma^1, v)$, we must have some positions $w \in V_2$ on the path $P(\sigma^2, v)$ such that $\sigma_2^2(w) \neq \sigma_2^1(w)$. Let us number all such positions as $w_1, \dots, w_k \in V_2$ for some $k \geq 1$, in the order we pass through them along the path $P(\sigma^2, v)$, when we start from v . Since σ is 1-basic, we cannot have $\sigma^2(w_k) \in V_T$, and thus we have $v_j := \sigma^2(w_j) \in V_1$ for all indices $j = 1, \dots, k$. Now, let us observe that the moves (v_j, w_j) exist, since this is an edge-symmetric terminal game, and thus they were considered in the uniform best improvement σ^1 , for all $j = 1, \dots, k$. Since they were not chosen, we must have the inequalities

$$\mathcal{L}^1(\sigma^1, v_j) \leq \mathcal{L}^1(\sigma^1, w_j)$$

for all $j = 1, \dots, k$ by Lemma 5 applied for the moves (v_j, w_j) with player $i = 1$, $j = 1, \dots, k$. We also have the equalities

$$\mathcal{L}^1(\sigma^1, v_j) = \mathcal{L}^1(\sigma^1, w_{j+1})$$

for $j = 1, \dots, k - 1$, due to our selection of vertices w_j , $j = 1, \dots, k$. These two groups of equations and inequalities imply that

$$\mathcal{L}^1(\sigma^1, w_1) \geq \mathcal{L}^1(\sigma^1, v_k).$$

Since we also have $\mathcal{L}^1(\sigma^2, v) = \mathcal{L}^1(\sigma^2, v_k) = \mathcal{L}^1(\sigma^1, v_k)$ and $\mathcal{L}^1(\sigma^1, v) = \mathcal{L}^1(\sigma^1, w_1)$, the stated inequalities follow. \square

Lemma 9 Assume that $\Gamma = (G, \mathcal{L})$ is a 2-person edge-symmetric terminal game satisfying conditions (CIW), conditions in Lemma 4, and that σ is a 1-basic situation in which all plays $P(\sigma, v)$ are finite. Assume further that σ^1 is obtained from σ by a uniform best improvement of player 1. Then situation σ^1 is 2-basic.

Proof Let us apply a uniform best improvement of player 2 to σ^1 , and denote the obtained situation by σ^2 .

Let us define $W = \{v \in V \mid N^+(v) \cap V_T \neq \emptyset\}$ and note that by Lemma 4 we have a terminal $t(v) \in V_T$ for all $v \in W$ such that $N^+(v) \cap V_T = \{t(v)\}$ holds.

Now, for a position $v \in W \cap V_2$ we have either $\sigma^2(v) = \sigma(v)$ and thus $\mathcal{L}^2(\sigma^2, v) = \mathcal{L}^2(\sigma, v)$, or we have $\sigma^2(v) \neq \sigma^1(v) = \sigma(v)$ and thus $\mathcal{L}^2(\sigma^2, v) < \mathcal{L}^2(\sigma^1, v) = \mathcal{L}^2(\sigma, v)$ by the definition of a uniform best improvement. In both cases the claimed property follows by our assumption that σ is 1-basic.

Finally, for positions $v \in W \cap V_1$ we have $\sigma^2(v) = \sigma^1(v)$. Thus, we either have $\sigma^2(v) = \sigma^1(v) = t(v)$, or $\sigma^1(v) \neq t(v)$, in which case we have $\mathcal{L}_1(\sigma^1, v) < \mathcal{L}^1(t(v))$ by the definition of a uniform best improvement. In this latter case we apply Lemma 8 and conclude $\mathcal{L}^1(\sigma^2, v) \leq \mathcal{L}^1(\sigma^1, v) < \mathcal{L}^1(t(v))$. \square

Proof of theorem 3 Let us first recall that the effective cost of an infinite play is defined as 0, and thus condition (CIW) is equivalent with saying that $\mathcal{L}^i(w) < 0$ for all players $i \in I$ and terminals $w \in V_T$. We also assume that the simplifying conditions in Lemma 4 hold.

Our proof idea is to show that if players alternate in making their uniform best improvements, then after a finite number of such iterations we arrive to a situation on which neither of the players can improve, showing that it is a UNE. While this idea would work starting with an arbitrary situation, our proof becomes simpler if we choose a special initial situation $\sigma^0 = (\sigma_1^0, \sigma_2^0)$, in which all plays are finite and which is 1-basic. According to Lemma 7 such a situation exists and is easy to construct.

After this, players, starting with player 1, alternate in making their uniform best improvements, as long as there is some change. This creates a series of situations, σ^j , $j = 1, 2, \dots$. To be very precise, situation σ^j for an odd $j \geq 1$ is obtained from σ^{j-1} by a uniform best improvement of player 1, while for an even $j \geq 2$ it is obtained from σ^{j-1} by a uniform best improvement of player 2.

Note that by our definitions, if a player has a uniform best improvement, it actually improves strictly his local costs in at least some of the positions it controls, and degrades in neither. If the above process terminates, the last situation is a UNE by definition.

On the other hand, a uniform best improvement by one of the players may actually increase some of the local costs for the other player, and thus the above procedure may start cycling through a number of situations, never terminating. To prove our theorem, we are going to show that such cycling cannot happen. To this end, we associate a quantity that can take only finitely many different values to every situation, and show that this decreases strictly in every step in the above process.

To a situation σ let us associate for $i \in I$

$$v^i(\sigma) = \sum_{v \in V_i} \mathcal{L}^i(\sigma, v),$$

and define the quantity associated to situations σ as

$$v(\sigma) = v^1(\sigma) + v^2(\sigma).$$

We claim that for all $j \geq 1$ we have

$$v(\sigma^{j+1}) < v(\sigma^j). \tag{4}$$

Note that by the definition of a uniform best improvement we have

$$\begin{aligned} v^1(\sigma^{j+1}) &< v^1(\sigma^j) && \text{if } j \text{ is even, and} \\ v^2(\sigma^{j+1}) &< v^2(\sigma^j) && \text{if } j \text{ is odd.} \end{aligned}$$

Thus, to prove our main claim (4), and consequently the theorem, it is enough to show that for all $j \geq 1$ and positions $v \in V$ we have

$$\mathcal{L}^{i(v)}(\sigma^{j+1}, v) \leq \mathcal{L}^{i(v)}(\sigma^j, v). \tag{5}$$

Note that since σ^0 is 1-basic, this claim follows by Lemma 8 for $j = 1$.

By applying Lemma 9, recursively, we can conclude that σ^j for an even $j \geq 2$ are all 1-basic, and for an odd $j \geq 1$ are all 2-basic.

Note also that inequality (5) holds for positions $v \in V_1$ whenever j is even, for $v \in V_2$ whenever j is odd, simply by the definition of a uniform best improvement. For the remaining cases, we can prove (5) by induction on j , starting with $j = 1$, and applying Lemma 8, with interchanging the players role's in it, alternately.

The above arguments prove (5), and thus (4) follows. Since $v(\sigma)$ can only take finitely many different values, this completes the proof of our theorem. \square

Remark 1 Though our paper is not of algorithmic nature, we would like to remark that such a UNE for an edge-symmetric terminal game Γ can be computed efficiently.

First of all, computing a uniform best improvement is equivalent with solving a so-called reachability problem in G , and thus can be done in $O(|E|)$ time.

A starting strategy σ^0 can be computed in $O(|V||E|)$ time, according to the procedure in Lemma 7.

According to our assumptions, as in Lemma 4 and to condition (CIW), we can assume w.l.o.g. that

$$-|V_T| \leq \mathcal{L}^i(w) < 0$$

for all terminals $w \in V_T$ and players $i \in I$. Thus we have

$$-|V||V_T| \leq v(\sigma) < 0$$

for all situations σ . Thus we need only to consider at most $|V||V_T|$ many uniform best improvements before termination, showing that our total complexity is $O(|V||V_T||E|)$.




Fig. 1 Terminal game Γ^2 from Boros et al. (2012). This example satisfies (TWO) and (SYM), but not (CIW) since **Player 1** prefers the infinite play, denoted by c to both terminals a_1 and a_2 . It is easy to verify that this game does not have a UNE




Fig. 2 A symmetrized variant of terminal game Γ^3 from Boros et al. (2012). This example satisfies (CIW) and (SYM), but not (TWO). It is easy to verify that this game does not have a UNE


Remark 2 We also mention that this theorem is sharp in the sense that all three listed conditions, (TWO), (SYM), and (CIW) are necessary, as we can demonstrate this by the next three examples:

- Consider Γ^2 from Boros et al. (2012) (see also Fig. 1). This example satisfies (TWO) and (SYM), but not (CIW) and has no UNE.
- Consider Γ^3 from Boros et al. (2012) and symmetrize it (see Fig. 2). This example satisfies (SYM) and (CIW), but has three players and has no UNE.
- Consider finally Γ^6 from Boros et al. (2012) (see also Fig. 3). This example satisfies (TWO) and (CIW), but not (SYM) and has no UNE.

Most of these claims are very easy to check, based on the graphical description of these terminal games given in Figs. 1, 2, and 3, except the last one, where the lack of a UNE is more complicated to check. We refer the reader to Boros et al. (2012) where this claim is demonstrated by describing the 8×8 normal form of Γ^6 .

6 NE-free examples

In this section we provide examples for edge-symmetric shortest path games that have no NE (and of course violate in some way the positivity condition). We also provide an example for a positive edge-symmetric shortest path game that has no UNE.



| Terminal Costs | | |
|----------------|----------|----------|
| | Player 1 | Player 2 |
| a_1 | -3 | -1 |
| a_2 | -4 | -5 |
| a_3 | -2 | -6 |
| a_4 | -1 | -3 |
| a_5 | -5 | -2 |
| a_6 | -6 | -4 |
| c | 0 | 0 |

Fig. 3 The terminal game Γ^6 from Boros et al. (2012). This example satisfies (CIW) and (TWO), but not (SYM). We refer the reader to Boros et al. (2012) for a proof of the claim that this game does not have a UNE

Let us recall (e.g., from Boros et al. 2017, 2018) that the effective cost of infinite plays can be defined in different ways, where differences are due to cycles in which the sum of edge lengths is zero (for some of the players). Accordingly, we provide two examples.

The first example shown in Fig. 4 is a 2-person nonzero sum shortest path game on an edge-symmetric digraph that has both positive and negative edge lengths and no cycle has length zero. The corresponding normal form is shown in Fig. 7 proving that this example has no NE.




Fig. 4 A 2-player edge-symmetric non-zero sum shortest path game example with no NE. Edge lengths are arbitrary real numbers, and there is no zero-cycle. Player RED controls the initial position s , and BLUE controls positions u and v . Both players are minimizing the length of the (infinite) path starting in the initial position s . Note that all stationary strategies yield a play that is ending in either a cycle yielding $\pm\infty$ as game value, or in a finite path terminating at t and yielding finite values for the players. For example, the play $s \rightarrow u \rightarrow s$ has a value $-\infty$ for RED, and ∞ for BLUE

The second example shown in Fig. 5 is a 2-person nonzero sum shortest path game on an edge-symmetric digraph that has nonnegative edge lengths but some of the cycles have zero length. In fact only cycles have zero length in which all edges have zero lengths. The corresponding normal form is shown in Fig. 8 proving that this example has no NE, either.

Let us add that all meaningful definitions of the cost of an infinite play (that may lead to a NE) agree on the following facts (see Boros et al. 2017, 2018): if for a situation σ the corresponding play $P = P(\sigma)$ ends in a cycle C and $\sum_{e \in C} \ell^i(e) > 0$, then $\ell^i(\sigma) = +\infty$, and if $\sum_{e \in C} \ell^i(e) < 0$, then $\ell^i(\sigma) = -\infty$. Furthermore, if all edges $e \in C$ have $\ell^i(e) = 0$, then $\ell^i(\sigma) = \sum_{e \in P \setminus C} \ell^i(e)$, or in other words, the effective cost of the play ending in cycle C is the length of the path of P leading to cycle C . In this sense, the above two examples show that positivity of the edge lengths in edge-symmetric shortest path games is “essential” to guarantee a NE.

Finally we show an example for a 2-person positive edge-symmetric shortest path game that has no UNE. This game, Γ^{6s} is shown in Fig. 6 is derived from the terminal game Γ^6 shown earlier. Note first that the counter clockwise moves have a large length for both players, larger than the length of any terminating move. This implies that in any UNE these moves will not be used. Let us also observe that from any position v and for any two terminals a_i and a_j , $i \neq j$ the $i(v)$ -lengths of the paths from v to these terminals (via the clockwise moves) compare exactly the same way as in the terminal game Γ^6 : the $i(v)$ -length of the $v \rightarrow a_i$ path is shorter than the $i(v)$ -length of the $v \rightarrow a_j$ path if and only if player $i(v)$ prefers terminal a_i to a_j in Γ^6 . Thus, any UNE in this shortest path game would correspond to a UNE in Γ^6 , and Boros et al. (2012) proved that no such UNE exist.




Fig. 5 A 2-player edge-symmetric non-zero sum shortest path game example with nonnegative edge lengths in which there is no NE. All edge lengths are nonnegative and some are equal to zero; there are full zero cycles. Player RED controls the initial position s , and BLUE controls positions u and v . Both players are minimizing the length of the (infinite) path starting at the initial position s . Note that all cycles are either positive (limit is $+\infty$), or have only arcs of length 0. For instance, the play $s \rightarrow v \rightarrow s$ has value 0 for RED and ∞ for BLUE




Fig. 6 A 2-person positive edge-symmetric shortest path game Γ^{6s} derived from the terminal game Γ^6 in Boros et al. (2012)

7 Open problems

Two important questions concerning the existence of NE remain open. The first one is about terminal games.

(Q1) Does an n -person terminal game satisfying condition (CIW) have a NE?

In this paper question **(Q1)** is answered in the positive for games on edge-symmetric digraphs. Moreover, in this case condition (CIW) is not essential.

It was shown in Boros et al. (2018) that in general (for non-symmetric digraphs) condition (CIW) is essential when $n > 2$, while for $n = 2$ (CIW) is not needed (Boros and Gurvich 2003). Interestingly, if the answer to **(Q1)** were negative then paradoxically there should be a terminal game satisfying (CIW) in which all NE are infinite, that is, realized by infinite plays; see (Boros et al. 2018).

The second open problem is about shortest path games.

(Q2) Does a positive 2-person shortest path game have a NE?

The answer is negative for more than two players, see (Gurvich and Oudalov 2014). In the present paper it is shown that the answer is negative for general

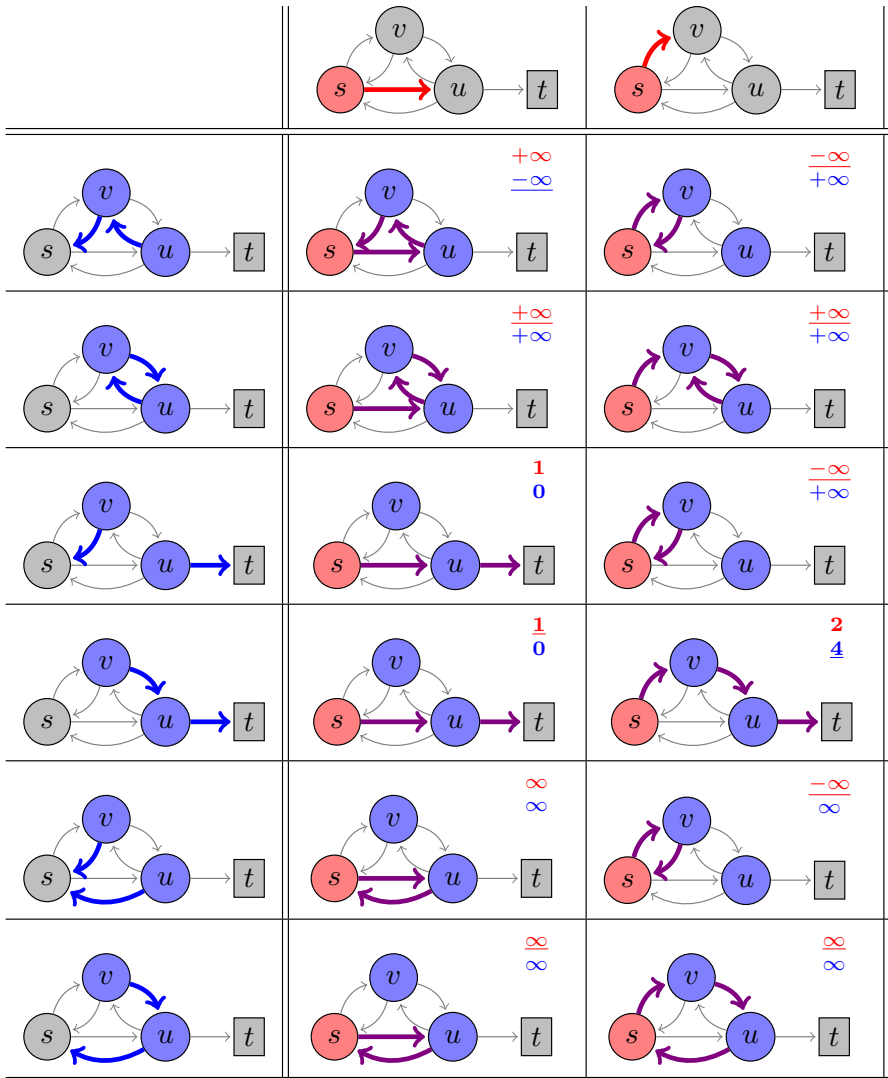


Fig. 7 Normal form of Example given in Fig. 4. The columns represent the two possible strategies of player RED and the rows represent the six possible strategies of player BLUE. Both players are minimizing their respective costs, indicated in the upper right corner in each cell. Row minimizing and column minimizing costs are underlined. Since no situation is simultaneously row and column minimizing, this example has no NE

edge-symmetric 2-person games, while it is positive for edge-symmetric positive games with any number of players.

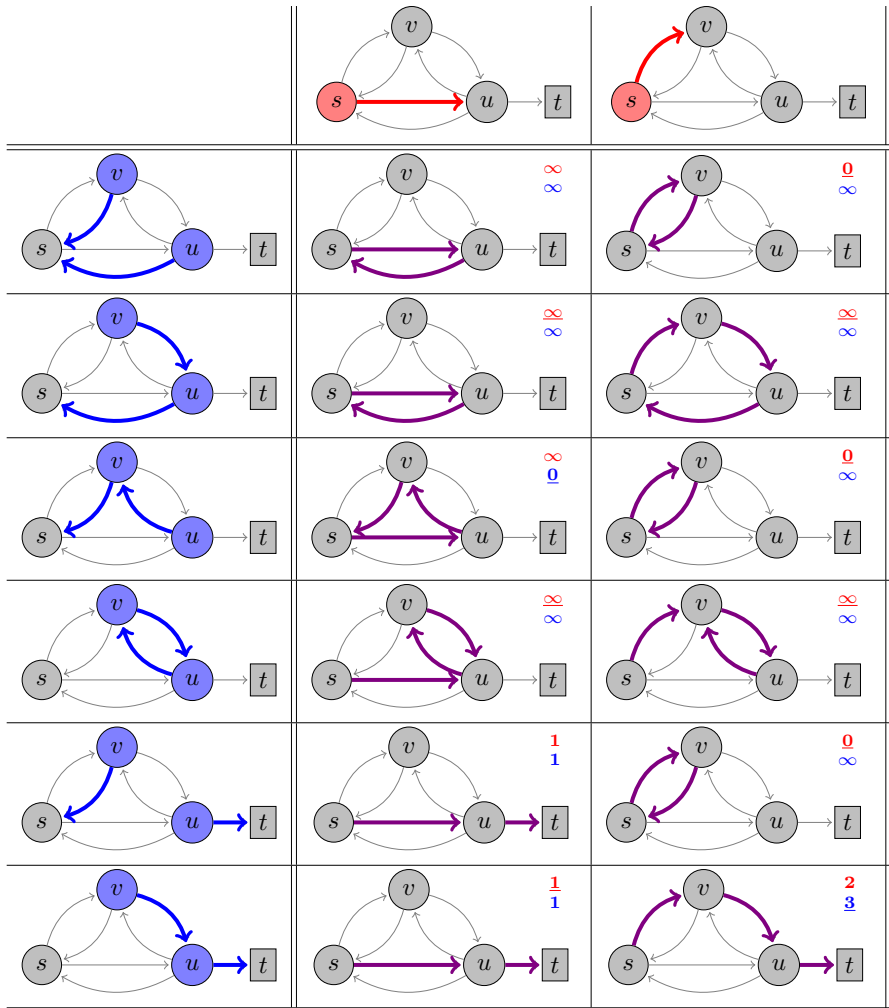


Fig. 8 Normal form of the example given in Fig. 5. The columns represent the two possible strategies of player RED and the rows represent the six possible strategies of player BLUE. Both players are minimizing their respective costs, indicated in the upper right corner in each cell. Since no situation is simultaneously row and column minimizing, this example has no NE

Acknowledgements The third and fourth authors were working within the framework of the HSE University Basic Research Program. The fourth author was supported in part by the state assignment topic no. 0063-2016-0003.

Funding Open access funding provided by Università degli Studi di Roma La Sapienza within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative

Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Boros E, Gurvich V (2003) On Nash-solvability in pure stationary strategies of positional games with perfect information which may have cycles. *Math Soc Sci* 46:207–241
- Boros E, Elbassioni K, Gurvich V, Makino K (2012) On Nash equilibria and improvement cycles in pure positional strategies for chess-like and Backgammon-like n -person games. *Discrete Math* 312(4):772–788
- Boros E, Elbassioni K, Gurvich V, Makino K (2017) A nested family of k -total effective rewards for positional games. *Int J Game Theory* 46(1):263–293
- Boros E, Elbassioni K, Gurvich V, Makino K (2018) Markov decision processes and stochastic games with total effective payoff. *Ann Oper Res*. <https://doi.org/10.1007/s10479-018-2898-8>
- Boros E, Gurvich V (2009) Why Chess and Backgammon can be solved in pure positional uniformly optimal strategies, RUTCOR Research Report, RRR-21-2009, Rutgers University
- Boros E, Gurvich V, Makino K, Wei S (2011) Nash-solvable two-person symmetric cycle game forms. *Discrete Appl Math* 159(15):1461–1487
- Boros E, Gurvich V, Milanič M, Oudalov V, Vičić J (2018) A three-person deterministic graphical game without Nash equilibria. *Discrete Appl Math* 243:21–38
- Everett H (1997) Recursive games, in M. Dresher, A. W. Tucker, and P. Wolfe (eds.) *Contributions to the Theory of Games, Annals of Mathematics Studies 3*, Princeton University Press. (1957) 67–78. Reprinted in H.W. Kuhn, ed., *Classics in Game Theory*, Princeton University Press
- Fraenkel AS, Scheinerman ER, Ullman D (1993) Undirected edge geography. *Theor Comput Sci* 112:371–381
- Gallai T (1958) Maximum-minimum Satze uber Graphen. *Acta Math Acad Sci Hungar* 9:395–434
- Gillette D (1957) Stochastic Games with Zero Stop Probabilities, In *Contributions to the Theory of Games, Vol. III, Annals of Mathematics Studies 39*:179–187
- Gurvich V (1975) Solution of positional games in pure strategies. *USSR Comput Math Math Phys* 15(2):74–87
- Gurvich V (1988) A stochastic game with complete information and without equilibrium situations in pure stationary strategies. *Russian Math Surv* 43(2):171–172
- Gurvich V (1989) Equilibrium in pure strategies. *Soviet Math Dokl* 38(3):597–602
- Gurvich V (2015) A four-person chess-like game without Nash equilibria in pure stationary strategies. *Bus Inf* 1(31):68–76
- Gurvich V (2018) Backward induction in presence of cycles. *Oxf J Log Comput* 28(7):1635–1646
- Gurvich V, Koshevoy G (2018) Monotone bargaining is Nash-solvable. *Discrete Appl Math* 250:1–15
- Gurvich V, Oudalov V (2014) On Nash-solvability in pure stationary strategies of the deterministic n -person games with perfect information and mean or total effective cost. *Discrete Appl Math* 167:131–143
- Lichtenstein D, Sipser M (1980) Go is polynomial-space hard. *J ACM* 27(2):393–401
- Schaefer TJ (1978) On the complexity of some two-person perfect-information games. *J Comput Syst Sci* 16(2):185–225
- Thuijsman F, Vrieze OJ (1987) The bad match, a total reward stochastic game. *Oper Res Spektrum* 9:93–99
- Thuijsman F, Vrieze OJ (1998) Total reward stochastic games and sensitive average reward strategies. *J Optim Theory Appl* 98:175–196
- Washburn AR (1990) Deterministic graphical games. *J Math Anal Appl* 153:84–96

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Endre Boros¹ · **Paolo Giulio Franciosa**²  · **Vladimir Gurvich**^{1,3} ·
Michael Vyalyi^{3,4,5}

✉ Paolo Giulio Franciosa
paolo.franciosa@uniroma1.it

Endre Boros
endre.boros@rutgers.edu

Vladimir Gurvich
vgurvich@hse.ru

Michael Vyalyi
vyalyi@gmail.com

¹ MSIS & RUTCOR, Business School, Rutgers University, 100 Rockafellar Road, Piscataway, NJ 08854, USA

² Department of Statistics, Sapienza University, P.le Aldo Moro 5, Rome 00185, Italy

³ Higher School of Economics (HSE), National Research University, Moscow, Russia

⁴ Federal Research Center “Computer Science and Control” of the Russian Academy of Sciences, Moscow, Russia

⁵ Department of Control and Applied Mathematics, Moscow Institute of Physics and Technology, Moscow, Russia