Random Zero-sum Dynamic Games on Infinite Directed Graphs

Luc Attia¹, Lyuben Lichev^{2, 3}, Dieter Mitsche^{2, 4}, Raimundo Saona^{5*}, Bruno Ziliotto^{1, 6}

 ¹Toulouse School of Economics, Université Toulouse Capitole, Institut de Mathématiques de Toulouse, CNRS UMR5219, Toulouse, France.
 ²University Jean Monnet, Saint-Etienne, France.
 ³Institute of Mathematics and Informatics, Bulgarian Academy of Science, Sofia, Bulgaria.
 ⁴Institute for Mathematical and Computational Engineering, Pontifícia Universidad Católica, Santiago, Chile.
 ⁵Institute of Science and Technology Austria, Klosterneuburg, Austria.
 ⁶CNRS, France.

*Corresponding author(s). E-mail(s): raimundo.saona@gmail.com; Contributing authors: luc.attia2@outlook.fr; lubetidobrilov@gmail.com; dmitsche@gmail.com; bruno.ziliotto@cnrs.fr;

Abstract

We consider random two-player zero-sum dynamic games with perfect information on a class of infinite directed graphs. Starting from a fixed vertex, the players take turns to move a token along the edges of the graph. Every vertex is assigned a payoff known in advance by both players. Every time the token visits a vertex, Player 2 pays Player 1 the corresponding payoff. We consider a distribution over such games by assigning i.i.d. payoffs to the vertices. On the one hand, for acyclic directed graphs of bounded degree and sub-exponential expansion, we show that, when the duration of the game tends to infinity, the value converges almost surely to a constant at an exponential rate dominated in terms of the expansion. On the other hand, for the infinite *d*-ary tree (that does not fall into the previous class of graphs), we show convergence at a double-exponential rate.

Keywords: Zero-sum games, Dynamic games, Random games, Directed graphs

MSC Classification: 91A05, 91A10, 91A25, 91A50, 05C05, 05C63

1 Introduction

Dynamic games on graphs

Zero-sum dynamic games with perfect information played on graphs [1] provide a powerful mathematical framework to analyze several important problems in mathematics and computer science. They correspond to stochastic games [2] with the restrictions that the transitions are deterministic and players play in turns. Considering a graph where every vertex has an outgoing edge, the game starts with a token at a vertex. Then, two players, who both know the position of the token at all times, alternate in choosing to move it along an outgoing edge to a neighboring vertex. Each vertex is assigned a uniformly bounded number called *payoff* and every time the token visits a vertex, Player 2 pays Player 1 its payoff. As the game evolves, the token moves from one vertex to another indefinitely, thus generating a sequence of payoffs.

Limit value

In the *n*-stage game, the objective of Player 1 is to maximize the mean payoff over nstages. Similarly, Player 2 aims to minimize this mean. The value of the n-stage game, denoted by v_n , is the maximum mean payoff Player 1 can guarantee irrespective of the actions of Player 2. Studying the behavior of dynamic games, or in general stochastic games, when the duration of the game tends to infinity, has been the subject of intense research over the last fifty years. One question that has received particular interest is the existence of a limit of the sequence $(v_n)_{n>1}$ as n tends to infinity. When it exists, this limit can be interpreted as the asymptotic mean payoff of Player 1 over an infinitely long game, and it thus stands out as a fundamental concept in the theory of stochastic games with long duration. A seminal result of Bewley and Kohlberg [3] states that, when the state space and the action sets are finite, the sequence $(v_n)_{n\geq 1}$ converges even for general stochastic games. This result has been extended in many directions, ranging from models with partial observation of the state and actions through models with unknown duration to models with other objective functions than mean payoff, see [4–7]. Proving convergence of $(v_n)_{n\geq 1}$ turns out to be a particularly delicate task when the state space is infinite. Indeed, positive results are scarce (see [8-10] for some recent advances) and counterexamples have been found [11].

Random dynamic games

Recently, a class of random zero-sum dynamic games on infinite graphs with vertices in \mathbb{Z}^d has been introduced under the name of *percolation games* [12]. In this model, as in usual dynamic games, each vertex is assigned a uniformly bounded payoff, and these values are known by the players before the game starts. The authors study a distribution over such games given by assigning i.i.d. random payoffs to the vertices. Then, the asymptotic behavior of the random value of the *n*-stage game, denoted by V_n , is studied. It is shown in [12] that, under the assumption that payoffs are bounded and every action increases the projection of the position of the token onto some axis, $(V_n)_{n\geq 1}$ converges almost surely to a deterministic limit value. Moreover, they provide exponential concentration estimates for $(V_n)_{n\geq 1}$ around the limit of their expectation. One important takeaway message of this result is that equipping the state space with a

particular graph structure (e.g., \mathbb{Z}^d) and assuming that payoffs have some probabilistic regularity (e.g., i.i.d. random variables) can ensure the existence of the limit of $(V_n)_{n\geq 1}$ even for an infinite graph.

Contributions

In this paper, we extend the results in [11] to a fairly more general class of graphs than \mathbb{Z}^d . In more detail, we introduce *directed games*, a class of games on acyclic directed graphs Γ where players move the token along the edges of Γ . On the one hand, under certain assumptions of weak transitivity and sub-exponential growth of Γ , we prove that $(V_n)_{n>1}$ is exponentially concentrated around a given deterministic limit value (so, in particular, $(V_n)_{n>1}$ a.s. converges to that value) and relate the convergence rate to the expansion of the graph (see Theorem 10). On the other hand, we consider the infinite d-ary tree where each vertex has exactly $d \geq 2$ children and every edge is directed from the parent to the child. These graphs do not belong to the previous class of games due to their exponential expansion away from the root. In this case, we show a stronger double-exponential concentration of the random variables $(V_n)_{n>1}$ around their respective expectations. In contrast to [12], the graphs that we consider here do not have to be transitive and, more importantly, they do not necessarily grow polynomially. This leads to several differences in the proof techniques. On a technical level, the novelty of our paper lies in treating a significantly more general class of graphs than the class of percolation games [12] while providing sharper concentration estimates. Note that the game being weakly transitive, and not simply transitive, introduces additional technicalities. Lastly, the proof of the second result on d-ary trees significantly differs from the arguments used in [12]: it is based on a pruning argument to avoid certain subtrees.

Related work

In our model, the payoffs associated with the states are the only random variables, and their realizations are known to both players before the game starts. Once the payoffs are revealed, the players play a deterministic game. Distributions over games, such as in our model of random dynamic games, have recently received growing attention. For example, studying random zero-sum games on graphs and their long-term behavior goes beyond the game theory community [13–16]. Indeed, on the one hand, a class of random games has been used to solve a well-known open problem regarding Probabilistic Finite Automata [17] (see [18] for an extension). On the other hand, the class of games studied in [12] has served as a toy model for contributing to a well-studied problem in the theory of PDEs called *stochastic homogenization* (see [12, Section 4] and [19, 20]).

Outline

This paper is organized as follows. In Section 2, we formally define our model and state our main results. In Section 3, we state some preliminary results for later use. Then, Sections 4 and 5 are dedicated to the proofs of our main results on controlled expansion graphs and d-ary trees, respectively. Finally, in Section 6, we show that our results cover all previously defined percolation games.

2 Model and main results

2.1 The model

A directed game is a dynamic game that consists of a locally finite directed graph Γ with infinite countable vertex set Z called the *state space*, an initial state $z_0 \in Z$ and a collection of independent and identically distributed (i.i.d.) random variables $(G_z)_{z \in Z}$ called *payoffs*. We assume that Γ has uniformly bounded degrees and contains neither directed cycles nor vertices with out-degree zero. The game is played by two players called Player 1 and Player 2. At the start of the game, the payoffs $(G_z)_{z \in Z}$ are sampled and presented to both players, who thus obtain perfect information. Then, a token is placed at the initial state z_0 . For every integer $i \geq 0$, given that the token is positioned at a state $z \in Z$ before stage i + 1, the following happens:

- if i is even, Player 1 moves the token to an out-neighbor z' of z in Γ ,
- if i is odd, Player 2 moves the token to an out-neighbor z' of z in Γ ,
- in all cases, Player 1 receives the payoff $G_{z'}$ from Player 2.

Note that, unlike the setting in [12], only one of the players performs a move at each stage (this is not a fundamental difference and makes our exposition clearer). We are mostly interested in the *n*-stage game consisting of the first n stages for (typically large) integers n.

A strategy of Player 1 (resp. Player 2) is a function $\sigma: \bigcup_{m\geq 0} Z^{2m+1} \to Z$ (respectively $\tau: \bigcup_{m\geq 0} Z^{2m+2} \to Z$) with the property that, for every $m \geq 0$ and $(z_0, z_1, \ldots, z_{2m+1}) \in Z^{2m+2}$, Γ contains the edge from z_{2m} to $\sigma(z_0, \ldots, z_{2m})$ (resp. from z_{2m+1} to $\tau(z_0, \ldots, z_{2m+1})$). We denote by Σ the collection of all strategies for Player 1 and by \mathcal{T} the collection of all strategies for Player 2.

Given a pair of strategies $(\sigma, \tau) \in \Sigma \times \mathcal{T}$, we define inductively the trajectory of the token by setting $z_{2i+1} \coloneqq \sigma(z_0, \ldots, z_{2i})$ and $z_{2i+2} \coloneqq \tau(z_0, \ldots, z_{2i+1})$ for every $i \ge 0$. This allows us to define the *n*-stage payoff function $\gamma_n^{z_0} \colon \Sigma \times \mathcal{T} \to \mathbb{R}$ by setting

$$\gamma_n^{z_0}(\sigma,\tau) \coloneqq \frac{1}{n} \sum_{i=1}^n G_{z_i}$$

Recall that the directed graph Γ is locally finite. Therefore, the *n*-stage game with initial state $z_0 = z \in Z$ is a perfect information finite game whose value is defined as usual by

$$V_n(z) := \max_{\sigma \in \Sigma} \min_{\tau \in \mathcal{T}} \gamma_n^z(\sigma, \tau) = \min_{\tau \in \mathcal{T}} \max_{\sigma \in \Sigma} \gamma_n^z(\sigma, \tau) \, ,$$

where the equality between maxmin and minmax is given by Kuhn's theorem [21, Theorem 1].

Moreover, we will say that a strategy $\sigma \in \Sigma$ (resp. $\tau \in \mathcal{T}$) is *optimal* for the *n*-stage game (starting from z) if σ maximizes $\min_{\tau \in \mathcal{T}} \gamma_n^z(\cdot, \tau)$ over Σ (resp. if τ minimizes $\max_{\sigma \in \Sigma} \gamma_n^z(\sigma, \cdot)$ over \mathcal{T}).

A classic question in the game-theoretic literature initiated by [3] is to ask for the convergence of the *n*-value as *n* grows to infinity. Since payoffs are random, V_n is a random variable. Therefore, we are interested in whether the sequence $(V_n)_{n\geq 1}$

converges a.s. to a constant. If no further assumptions are imposed, $(V_n)_{n\geq 1}$ does not

necessarily converge, as the following example shows. **Example 1.** For all integers $m \ge 0$, set $n_m \coloneqq 2^{2^{2m}}$ and $n'_m \coloneqq 2^{2^{2m+1}}$ and consider the case where Γ is a directed tree (all edges being directed away from the root) where each node of even height has only one child, while each node with odd height k has two children if k = 1 or $k \in [n_m, n'_m)$ for some $m \ge 0$, and it has only one child if $k \in [n'_m, n_{m+1})$. Moreover, let the payoffs be i.i.d. Bernoulli random variables with parameter 1/2. In particular, for every $m \geq 1$, in the n_m -stage game, Player 2 has only one choice most of the time, while in the n'_m -stage game, she has two choices most of the time. Since Player 2 can pick a vertex with payoff 0 (if it is present), and pick an available vertex otherwise, one can show that a.s.

$$\limsup_{m \to \infty} V_{n'_m} \le \frac{3}{8} < \frac{1}{2} = \lim_{m \to \infty} V_{n_m} \,.$$

Indeed, while Player 1 never has a choice in the n'_m -game (implying that the mean payoff over the odd states visited by the token a.s. converges to 1/2), Player 2 can ensure with the above strategy that the mean payoff over the even states visited by the token a.s. converges to 1/4, which yields that a.s. $\limsup_{m\to\infty} V_{n'_m} \leq 3/8$. At the same time, for every $\varepsilon > 0$, Chernoff's bound for the Binomial distribution $Bin(n_m, 1/2)$ and a union bound over the $O(2^{n'_{m-1}})$ vertices at level n_m in Γ shows that V_{n_m} is in the interval $[1/2 - \varepsilon, 1/2 + \varepsilon]$ with probability very close to 1. In particular, a.s. $(V_n)_{n\geq 1}$ does not converge. Therefore, to ensure convergence, we will need further structural assumptions on the graph.

Before turning to our results, we provide some vocabulary. Given a vertex $z \in Z$, a descendant of z (in Γ) is a vertex that can be reached from z by a directed path in Γ . We say that z and z' are equivalent if the two subgraphs of Γ induced by the descendants of z and by the descendants of z', respectively, are isomorphic (as directed graphs).

Definition 2. The graph Γ is weakly transitive if there is a state z^* and an integer $M \geq 0$ such that the following holds: for each state $z \in Z$, in the game with initial state $z_0 = z$, each player has a strategy that, independently of the moves of the opponent, ensures that the token is placed at a state equivalent to z^* after an even number of $\ell \leq M$ stages.

Note that all vertex-transitive graphs are weakly transitive with M = 0. We also remark that the parity of ℓ is important to ensure that, when the token reaches z^* , it is Player 1's turn to make a move. For similar notions of generalized vertex-transitive graphs, see for example [22].

In the remainder of the paper, we always assume that Γ is weakly transitive. The next two subsections present two types of directed games used in our main results.

2.1.1 Weakly transitive games with sub-exponential expansion

One of the main difficulties in the analysis of the asymptotic behavior of $(V_n)_{n\geq 1}$ is to make use of the independence of the payoffs $(G_z)_z$. To this end, we use the partial order introduced by the directed graph Γ on the state space. More formally, let $z_1, z_2 \in \mathbb{Z}$. We say that $z_1 \leq z_2$ if z_2 is a descendant of z_1 . Then, for all $z \in Z$ and $n \geq 1$, denote $Z_n(z)$ the set of *reachable states* from z in exactly n steps, and $Z^{(n)}(z)$ the ones reachable in at most n steps. Since Γ is locally finite, $Z^{(n)}(z)$ is a finite partially ordered set. Denote h(n, z) its height, that is, the longest path in $Z^{(n)}(z)$. Note that, starting from z, after h(n, z) stages, the game never returns to a state in $Z^{(n)}(z)$ as Γ is acyclic, thus separating $Z^{(n)}(z)$ from the future of the game. Hence, h(n, z) can be considered as the waiting time after which the future of the game becomes independent of $Z^{(n)}(z)$. In this regard, we define the largest possible waiting time, which we call *transient speed function*, as

$$h: n \in \mathbb{N} \mapsto \sup_{z \in Z} h(n, z) \in \mathbb{N} \cup \{\infty\}.$$

We note that Mirsky's theorem (see [23, Theorem 2]) gives us a dual interpretation of h(n, z) as the minimum number of antichains needed to partition $Z^{(n)}(z)$. Thus, h(n) may be seen as an upper bound on these minima. Following this point of view, we expect to obtain a concentration inequality parameterized by h(n) since the payoffs are i.i.d. In our analysis, this concentration has to be sufficiently strong to overcome a union bound over all possible states from which the game may continue, which is quantified by the growth of $Z^{(n)}(z)$. Formally, we will handle concentration inequalities of the form $\mathbb{P}(|X - \mathbb{E}[X]| \ge t)$, so we define the following key function $\psi \colon \mathbb{N} \times (0, \infty) \to \mathbb{R}$ by

$$\psi(n,t) \coloneqq \exp\left(-\frac{t^2 n^2}{2h(n)}\right) \max_{z \in \mathbb{Z}} |Z^{(2n)}(z)|.$$

Our main goal is to analyze directed games in which the size of the sets $Z^{(n)}(z)$ does not grow too fast as n grows to infinity, which we formalize as follows.

Definition 3 (δ -transient games). For a fixed $\delta > 0$, a directed game on a graph Γ with vertex set Z is called δ -transient if there exists a sequence $(\varepsilon_n)_{n\geq 1}$ of positive real numbers such that $\varepsilon_n + \psi(n, \varepsilon_n) = O(n^{-\delta})$. Note that, unless explicitly mentioned otherwise, our asymptotic notation will be taken with respect to $n \to \infty$.

Remark 4. The concept of δ -transient games is only relevant for $\delta \in (0, 1/2)$. Indeed, Definition 3 requires that $(\psi(n, \varepsilon_n))_{n\geq 1}$ converges to zero. Therefore, since $h(n) \geq n$, this implies that $n = o(\varepsilon_n^2 n^2)$, so $\varepsilon_n^2 n \to \infty$.

Remark 5. A sufficient condition under which a directed game is δ -transient is the following: there exist real numbers $\alpha \in [0, 2 - 2\delta)$ and $\beta \in [0, 2 - 2\delta - \alpha)$ such that $h(n) = O(n^{\alpha})$ and $\max_{z \in \mathbb{Z}} |Z^{(n)}(z)| = \exp(O(n^{\beta})).$

Note that the definition of a δ -transient game is independent of the payoffs and only makes assumptions on the state space. In Section 2.1.1, we give a few examples of δ -transient games.

Oriented games. Oriented games were introduced by Garnier and Ziliotto [12] and are defined as follows. Fix an integer $d \ge 1$, and denote by e_i the *d*-dimensional vector with 1 in coordinate *i* and 0 in all other d - 1 coordinates. Given positive integers $n_1, \ldots, n_d \ge 1$, a directed graph Γ with vertex set $Z \subseteq \mathbb{Z}^d$ is called (n_1, \ldots, n_d) *invariant* (or simply *invariant*) if, for every $i \in [1, d]$, the translation at vector $n_i e_i$ is a graph isomorphism for Γ . A directed weakly transitive game is called *oriented* if

its underlying graph Γ is invariant and there exists $u \in \mathbb{R}^d \setminus \{0\}$ such that, for every directed edge (z, w) in Γ , we have $(w - z) \cdot u > 0$ (here, \cdot denotes the usual scalar product of vectors in \mathbb{R}^d). We show the following proposition.

Proposition 6. Every oriented game is δ -transient for all $\delta \in (0, 1/2)$.

The following two classes of games present particular examples of oriented games. **Example 7** (Games on tilings, see Figure 1). A tiling is a periodic partition of the plane into translations of one or several polygonal shapes, called tiles, with vertices in \mathbb{Z}^2 . Tilings naturally define planar graphs whose vertex set coincides with the corners of the tiles and two vertices are connected by an edge if these can be connected by following the boundary of a tile without meeting another vertex on the way. Equipping the edges of this graph with suitable orientations defines an oriented game.



Fig. 1 The figure depicts part of a tiling with two types of square tiles. The vertices and the edges of the planar graph originating from the tiling are depicted in blue and red, respectively. Each horizontal edge is oriented from left to right and every vertical edge is oriented from bottom to top. One may choose z^* to be the bottom left vertex of a small square and M = 6.

Example 8 (Games on directed chains of graphs). Fix a finite vertex-transitive graph H with vertex set V(H) and edge set E(H), and a bi-infinite sequence of copies $(H_i)_{i \in \mathbb{Z}}$ of H. For every $i \in \mathbb{Z}$ and $u \in V(H)$, denote by u_i the vertex in H_i corresponding to u. We call an H-chain the graph Γ_H with vertices $\bigcup_{i \in \mathbb{Z}} V(H_i)$ and edges $\{u_i v_{i+1} : i \in \mathbb{Z}, uv \in E(H)\}$.

Games on H-chains can be seen as instances of oriented games on \mathbb{Z} . Indeed, fixing h = |V(H)|, one may identify the vertices of H_i with the integers in the interval [ih + 1, (i + 1)h] for all $i \in \mathbb{Z}$ in a translation-invariant way.

Weakly transitive games with controlled expansion. In the following paragraphs, we show that one can construct a tree T with an arbitrary growth that is faster than linear but slower than exponential. In particular, we have the following result.

Proposition 9. For every $\delta \in (0, 1/2)$, there exist games that are δ -transient but are not δ' -transient for every $\delta' > \delta$.

Fix an arbitrary infinite rooted tree T with root r and a family of vertex-disjoint infinite paths $(P_v)_{v \in V(T)}$ where the path P_v intersects T at a unique vertex v. Define $\Gamma = T \cup (\bigcup_{v \in V(T)} P_v)$ as the tree rooted in r and with all edges oriented away from r. Thus, Γ is a directed rooted tree larger than T. Since a single move of each player is sufficient to place the token at the second vertex of some infinite path among $(P_v)_{v \in V(T)}$, the game is weakly transitive.

Recall that $Z_n(z)$ consists of all vertices at distance n from z. Therefore, the sequence $(Z_n(r))_{n\geq 1}$ is a partition of the vertex set Z of Γ and $Z_n(r)$ can only be visited once. Since $Z^{(n)}(r)$ can be partitioned into n antichains Z_1, Z_2, \ldots, Z_n , we have that $h(n, r) \leq n$.

Let us show that we can control the growth speed of $\max_{z \in Z} |Z^{(2n)}(z)|$. Consider a set of non-negative integers $L = \{\ell_i : i \ge 1\}$ with $\ell_1 < \ell_2 < \ldots$ and let every vertex of T in level ℓ have two children if $\ell \in L$ and one child otherwise. Moreover, suppose that $\ell_1 = 0$ and $(\ell_i - \ell_{i-1})_{i\ge 1}$ is a non-decreasing sequence. Then, one can readily check that, for every $n \ge 1$, $\max_{z \in Z} |Z^{(n)}(z)| = |Z^{(n)}(r)|$, and h(n, z) = h(n, r) = n. Indeed, for every $k, n \ge 1$ and a vertex $z \in Z$ on level k, using the assumptions that $\ell_1 = 0$ and $(\ell_i - \ell_{i-1})_{i\ge 1}$ is a non-decreasing sequence, we get

$$|Z^{(n)}(z) \setminus Z^{(n-1)}(z)| = 2^{|L \cap \{k,\dots,k+n-1\}|} \le 2^{|L \cap \{0,\dots,n-1\}|} = |Z^{(n)}(r) \setminus Z^{(n-1)}(r)|.$$

Thus, for every integer $n \ge 0$, $|Z^{(n)}(r)| = 1 + \sum_{i=0}^{n-1} 2^{|L \cap \{0, \dots, i\}|}$.

2.1.2 Directed games on *d*-ary trees

We turn our attention to a natural example of a directed game where the set of reachable states after n steps grows exponentially with n. Note that, for all $\delta > 0$, it is not a δ -transient game. Fix an integer $d \ge 2$ and let T be an *infinite d-ary tree*, that is, a tree where every vertex has d children, with vertex set Z where every edge is directed from the parent to the child. We fix an arbitrary initial vertex z_0 and, for every integer $i \ge 0$, we define Z_i to be the set of vertices in Z that can be reached from z_0 by exactly i steps and also denote $Z_{\text{even}} \coloneqq \bigcup_{i\ge 0} Z_{2i}$ and $Z_{\text{odd}} \coloneqq \bigcup_{i\ge 0} Z_{2i+1}$. Note that, for every $n \ge 1$, the random variables $(V_n(z))_{z\in Z}$ have the same distribution. Hence, in this setting, we often omit the dependence of V_n in z.

2.2 Main results

Our first main result shows sharp concentration for the *n*-value of δ -transient games around a deterministic constant.

Theorem 10. Fix $\delta \in (0, 1/2)$. Consider a δ -transient directed game with transient speed h and i.i.d. payoffs $(G_z)_{z \in Z}$ supported on the interval [0, 1]. Then, there exist constants $v_{\infty} \in [0, 1]$ and K > 0 such that, for all $n \ge 1, t \ge 0$, and $z \in Z$,

$$\mathbb{P}\left(|V_n(z) - v_{\infty}| \ge t + Kn^{-\delta}\right) \le 2\exp\left(-\frac{t^2n^2}{2h(n)}\right) \,.$$

Consequently, for all $z \in Z$, $(V_n(z))_{n \geq 1}$ converges almost surely to v_{∞} .

Our second main result shows that the *n*-value of the directed game on a *d*-ary tree is tightly concentrated around a constant.

Theorem 11. Fix an integer $d \ge 2$. Consider a directed game on the d-ary tree with *i.i.d.* payoffs supported on the interval [0, 1]. Then, there exists a constant $v_{\infty} \in [0, 1]$ such that, for every $\delta \in (0, 1/2)$, there exists K > 0 such that, for all $n \ge 1$, $t \ge 0$, and $z \in Z$,

$$\mathbb{P}(|V_n(z) - v_{\infty}| \ge t + 2t^2 + Kn^{-\delta}) \le \exp\left(-\frac{1}{6}\exp\left(\frac{t^2n}{4}\right)\right) \,.$$

Consequently, for all $z \in Z$, $(V_n(z))_{n \ge 1}$ converges almost surely to v_{∞} .

On a high level, the proofs of both theorems contain two main steps.

- 1. The first step relies on concentration arguments showing that V_n is close to $\mathbb{E}[V_n]$ with high probability. While standard concentration tools are sufficient for our proof of Theorem 10, the stronger probabilistic bound in Theorem 11 requires an additional boosting obtained by dividing the first n levels of the d-ary tree into two groups of consecutive levels and treating the n-stage game as two consecutive games on k and n k stages, respectively.
- 2. The second step uses the structure of the underlying graph to show that $(\mathbb{E}[V_n])_{n\geq 1}$ satisfies a certain subadditivity assumption, which allows us to conclude that $(\mathbb{E}[V_n])_{n\geq 1}$ converges to a constant v_{∞} , and moreover, $|\mathbb{E}[V_n] v_{\infty}|$ is polynomially small.

The proof of Proposition 6 relies on a simple explicit construction.

Perspectives.

The proofs of Theorems 10 and 11 have a similar structure but use different arguments. A challenging research question would be to prove convergence of $(V_n)_{n\geq 1}$ and concentration bounds in all weakly transitive directed games, irrespective of the expansion speed of the underlying graph, thus unifying Theorems 10 and 11.

3 Preliminary results

In our proofs, we make use of the well-known *bounded difference inequality*, also known as *McDiarmid's inequality*.

Lemma 12 ([24, Corollary 2.27]). Fix a function $f: \Lambda_1 \times \cdots \times \Lambda_N \to \mathbb{R}$ and let Y_1, \ldots, Y_N be independent random variables taking values in $\Lambda_1, \ldots, \Lambda_N$, respectively. Suppose that there are positive constants c_1, \ldots, c_N such that, for every two vectors

 $z, w \in \Lambda_1 \times \cdots \times \Lambda_N$ that differ only in the k-th coordinate, we have $|f(z) - f(w)| \le c_k$. Then, for every $t \ge 0$, the random variable $X = f(Y_1, \ldots, Y_N)$ satisfies

$$\mathbb{P}(X - \mathbb{E}[X] \ge t) \le \exp\left(-\frac{t^2}{2\sum_{i=1}^N c_i^2}\right).$$
$$\mathbb{P}(X - \mathbb{E}[X] \le -t) \le \exp\left(-\frac{t^2}{2\sum_{i=1}^N c_i^2}\right).$$

We also use the following result that states convergence of almost subadditive sequences.

Lemma 13 ([25, Theorem 23]). Fix an increasing function $\phi \colon \mathbb{N} \to (0, \infty)$ such that the sum of $(\phi(n)/n^2)_{n\geq 1}$ is finite, and a function $f \colon \mathbb{N} \to \mathbb{R}$ such that, for all $n \in \mathbb{N}$ and all integers $m \in [n/2, 2n]$, $f(n+m) \leq f(n) + f(m) + \phi(n+m)$. Then, there exists $\ell \in \mathbb{R} \cup \{-\infty\}$ such that

$$\left(\frac{f(n)}{n}\right) \xrightarrow[n \to \infty]{} \ell \,.$$

4 δ -transient games: proof of Theorem 10

To begin with, by using a suitable partition of the state space into subsets that are visited at most once, we show that the value of the n-game is well concentrated around its expected value. Note that the next lemma holds for weakly transitive games in general and will be reused in Section 5.

Lemma 14. For every $z_0 \in Z$, $n \ge 1$, and $t \ge 0$,

$$\mathbb{P}(V_n(z_0) - \mathbb{E}[V_n(z_0)] \ge t) \le \exp\left(-\frac{t^2 n^2}{2h(n)}\right),$$
$$\mathbb{P}(V_n(z_0) - \mathbb{E}[V_n(z_0)] \le -t) \le \exp\left(-\frac{t^2 n^2}{2h(n)}\right)$$

Proof. Let us fix $z_0 \in Z$ and abbreviate $V_n = V_n(z_0)$, $Z_n = Z_n(z_0)$ and $Z^{(n)} = Z^{(n)}(z_0)$. Define the (random) vectors $X_k = (G_z)_{z \in Z_k} \in [0,1]^{|Z_k(z_0)|}$. Then, since $Z^{(n)} \subseteq Z_{[h(n)]}$, V_n can be written as $f(X_1, \ldots, X_{h(n)})$ for some function $f: [0,1]^{|Z_1|} \times \cdots \times [0,1]^{|Z_{h(n)}|} \to \mathbb{R}$. Moreover, for every integer $k \in [1, h(n)]$, the token visits the set Z_k at most once and therefore, for every pair of strategies $(\sigma, \tau) \in \Sigma \times \mathcal{T}, \gamma_n^{z_0}(\sigma, \tau)$ varies by at most 1/n as a function of X_k . Hence, for every choice of vectors $(x_i)_{i=1}^{h(n)} \in [0,1]^{|Z_1|} \times \cdots \times [0,1]^{|Z_{h(n)}|}$ and $x'_k \in [0,1]^{|Z_k|}$,

$$|f(x_1, \dots, x_k, \dots, x_{h(n)}) - f(x_1, \dots, x'_k, \dots, x_{h(n)})| \le \frac{1}{n}.$$

Lemma 12 applied to V_n finishes the proof.

In the remainder of the proof, we show that the expected n-value converges to a constant polynomially fast. Next, we show how to control the difference of the values of games of different lengths.

Lemma 15. Fix integers $n \ge 1$ and $k \in [1, n]$. Then, for every $z_0 \in Z$, $|V_n(z_0) - V_{n-k}(z_0)| \le k/n$.

Proof. Observe that $nV_n(z_0) \ge (n-k)V_{n-k}(z_0)$ and $nV_n(z_0) \le (n-k)V_{n-k}(z_0) + k$. Indeed, in the *n*-stage game, Player 1 (respectively Player 2) may first play according to an optimal strategy for the (n-k)-stage game, and play arbitrarily during the last k stages. Hence, $|n(V_n(z_0) - V_{n-k}(z_0))| \le \max(kV_{n-k}(z_0), k - kV_{n-k}(z_0)) \le k$, which implies the statement of the lemma.

The next lemma shows that starting from different initial states changes the expected n-value only slightly when n is large.

Lemma 16. For every $z_0 \in Z$, $|\mathbb{E}[V_n(z_0)] - \mathbb{E}[V_n(z^*)]| = O(n^{-\delta})$.

The main ideas of the proof are as follows: Given states z, z^* , the assumption of weak transitivity allows us to bound from below (respectively from above) $V_n(z)$ using the minimum (respectively the maximum) of the *n*-values over the nearby states equivalent to z^* . Moreover, by δ -transience, the graph Γ does not expand too quickly, which implies that, by a union bound, the expectation of the minimum (respectively the maximum) of the said *n*-values is approximately equal to the minimum (respectively the maximum) of their expectations.

Proof of Lemma 16. Fix an initial state $z_0 \in Z$ and denote by E the set of states in $Z^{(M)}(z_0)$ that are equivalent to z^* . By Definition 2, $E \neq \emptyset$ and, independently of the moves of Player 2, and Player 1 can ensure that the token is at a state in E after an even number of $\ell \leq M$ stages. Hence, using Lemma 15, we have

$$nV_n(z_0) \ge (n-M)\min_{z \in E} V_{n-M}(z) \ge (n-M)\min_{z \in E} (V_n(z) - M/n) \ge \min_{z \in E} nV_n(z) - 2M.$$
(1)

Now, we bound from below the expectation of the right-hand side. Combining Lemma 14 and the choice of ε_n from Definition 3, we have

$$\mathbb{E}\left[\min_{z\in E} V_n(z)\right] \ge (\mathbb{E}[V_n(z^*)] - \varepsilon_n)(1 - \mathbb{P}(\exists z\in E: V_n(z) \le \mathbb{E}[V_n(z)] - \varepsilon_n))$$

$$\ge (\mathbb{E}[V_n(z^*)] - \varepsilon_n)\left(1 - \exp\left(-\frac{\varepsilon_n^2 n^2}{2h(n)}\right) \max_{z\in Z} |Z^{(M)}(z)|\right) \qquad (2)$$

$$\ge (\mathbb{E}[V_n(z^*)] - \varepsilon_n)(1 - \psi(n,\varepsilon_n)) = \mathbb{E}[V_n(z^*)] - O(n^{-\delta}),$$

where the second inequality comes from a union bound, the third one uses that $Z^{(M)}(z) \subseteq Z^{(2n)}(z)$ for every $z \in Z$, and the equality is implied by the fact that $\varepsilon_n + \psi(n, \varepsilon_n) = O(n^{-\delta})$. Thus, taking expectations on both sides of (1) and using (2) shows that

$$\mathbb{E}[V_n(z_0)] \ge \mathbb{E}[V_n(z^*)] - O(n^{-\delta} + 2M/n) = \mathbb{E}[V_n(z^*)] - O(n^{-\delta}).$$
(3)

Similarly, Player 2 can ensure that the token reaches a state in E after an even number of $\ell \leq M$ stages. Hence,

$$nV_n(z_0) \le (n-M) \max_{z \in E} V_{n-M}(z) + M \le \max_{z \in E} nV_n(z) + M.$$
(4)

At the same time, similarly to (2), $\mathbb{E}\left[\max_{z \in E} V_n(z)\right]$ is bounded from above by

$$(\mathbb{E}[V_n(z^*)] + \varepsilon_n)(1 - \mathbb{P}(\exists z \in E : V_n(z) \ge \mathbb{E}[V_n(z)] + \varepsilon_n)) + \mathbb{P}(\exists z \in E : V_n(z) \ge \mathbb{E}[V_n(z)] + \varepsilon_n),$$

which is at most $\mathbb{E}[V_n(z^*)] + (\varepsilon_n + \psi(n, \varepsilon_n)) = \mathbb{E}[V_n(z^*)] + O(n^{-\delta})$. Combining this with (4) shows that $\mathbb{E}[V_n(z_0)] \leq \mathbb{E}[V_n(z^*)] + O(n^{-\delta})$, and, together with the lower bound in (3), this finishes the proof.

Next, we show that the expected value of V_n converges as $n \to \infty$. Lemma 17. There is a constant v_{∞} such that, $|\mathbb{E}[V_n(z)] - v_{\infty}| = O(n^{-\delta})$, for every $z \in \mathbb{Z}$.

As before, we start with an outline of the proof. We first use Lemma 16 to show that the initial state of the game is not important. Then, we justify that the sequence $(-n\mathbb{E}[V_n(z)])_{n\geq 1}$ is approximately subadditive, and use Lemma 13. The subadditivity is based on the simple observation that, for every player, playing optimally in the (m+n)-game is better than playing optimally in the initial *m*-stage game (based only on the *m*-th neighborhood of the initial state), and then do the same in the subsequent *n*-stage game.

Proof of Lemma 17. By Lemma 16, it is sufficient to show the lemma when $z_0 = z^*$; we abbreviate $V_n = V_n(z^*)$ and $Z^{(n)} = Z^{(n)}(z^*)$ for convenience. First, we show that $\mathbb{E}[V_n]$ converges to a limit $v_{\infty} \in \mathbb{R}$ as $n \to \infty$. By Lemma 14 and a union bound, for all $t \ge 0$,

$$\mathbb{P}(\exists z \in Z^{(2n)} : |V_n(z) - \mathbb{E}[V_n(z)]| \ge t) \le \sum_{z \in Z^{(2n)}} \mathbb{P}(|V_n(z) - \mathbb{E}[V_n(z)]| \ge t) \\
\le 2 \exp\left(-\frac{t^2 n^2}{2h(n)}\right) \max_{z \in Z} |Z^{(2n)}(z)| = 2\psi(n, t),$$
(5)

where the factor of 2 in the last inequality appears since we bound both the upper and the lower tail of $V_n(z)$.

By definition of δ -transient game, there exists $(\varepsilon_n)_{n\geq 1}$ such that $\varepsilon_n + \psi(n, \varepsilon_n) = O(n^{-\delta})$. Denote by E the set of vertices in $Z^{(2n)}$ that are equivalent to z^* . Now, Lemma 16 implies that there is a constant K' > 0 such that, for every $z \in Z$ and $n \geq 1$, $|\mathbb{E}[V_n(z)] - \mathbb{E}[V_n]| \leq K' n^{-\delta}$. Combining this with (5), we get that

$$\mathbb{P}\left(\min_{z\in Z^{(2n)}} V_n(z) \le \mathbb{E}[V_n] - \varepsilon_n - K'n^{-\delta}\right) \le \mathbb{P}\left(\min_{z\in Z^{(2n)}} |V_n(z) - \mathbb{E}[V_n(z)]| \ge \varepsilon_n\right)$$

$$\leq \mathbb{P}(\exists z \in E, |V_n(z) - \mathbb{E}[V_n(z)]| \geq \varepsilon_n) \\ \leq 2\psi(n, \varepsilon_n) = O(n^{-\delta}).$$

In particular, it follows directly that

$$\begin{split} \mathbb{E}\left[\min_{z\in Z^{(2n)}}V_n(z)\right] &\geq \left(\mathbb{E}[V_n] - \varepsilon_n - K'n^{-\delta}\right) \,\mathbb{P}\left(\min_{z\in Z^{(2n)}}V_n(z) \geq \mathbb{E}[V_n] - \varepsilon_n - K'n^{-\delta}\right) \\ &\geq \left(\mathbb{E}[V_n] - \varepsilon_n - K'n^{-\delta}\right) \,\left(1 - 2\psi(n,\varepsilon_n)\right) \\ &\geq \mathbb{E}[V_n] - 2(\psi(n,\varepsilon_n) + \varepsilon_n) - K'n^{-\delta} \,. \end{split}$$

Now, fix an integer $m \in [1, 2n]$ and consider the (m + n)-stage game. Suppose that Player 1 plays according to an optimal strategy for the *m*-stage game up to stage *m* and, once the *m*-stage game terminates at a state z_m , continues to play according to an optimal strategy for the subsequent *n*-stage game. Note that $z_m \in Z^{(2n)}$, so the above strategy of Player 1 for the first m + n steps guarantees a gain of $\frac{m}{m+n}V_m + \frac{n}{m+n}\min_{z \in Z^{(2n)}}V_n(z)$. Thus,

$$(m+n)\mathbb{E}[V_{m+n}] \ge m\mathbb{E}[V_m] + n\mathbb{E}\left[\min_{z\in Z^{(2n)}} V_n(z)\right]$$

$$\ge m\mathbb{E}[V_m] + n\mathbb{E}[V_n] - 2n(\psi(n,\varepsilon_n) + \varepsilon_n) - K'n^{1-\delta}.$$
(6)

Since $\psi(n, \varepsilon_n) + \varepsilon_n = O(n^{-\delta})$, there is a constant K'' > 0 such that, for all $n \ge 1$,

$$2n(\psi(n,\varepsilon_n)+\varepsilon_n)+K'n^{1-\delta}\leq 2K''n^{1-\delta}.$$

Thus, using Lemma 13 with $f: n \mapsto -n\mathbb{E}[V_n]$ and $\phi: n \mapsto 2K''n^{1-\delta}$ (note that ϕ is increasing and $\sum_{n\geq 1} \phi(n)/n^2 = 2K'' \sum_{n\geq 1} 1/n^{1+\delta} < \infty$) implies that $\mathbb{E}[V_n]$ converges to a limit $v_{\infty} \in \mathbb{R} \cup \{\infty\}$ as $n \to \infty$. Note that v_{∞} is in [0, 1] since this is the support of all payoff variables.

Finally, using (6) with m = n, for every $n \ge 1$, we have that

$$\mathbb{E}[V_{2n}] \ge \mathbb{E}[V_n] - (\psi(n,\varepsilon_n) + \varepsilon_n) - \frac{K'n^{-\delta}}{2} \ge \mathbb{E}[V_n] - K''n^{-\delta}.$$

In particular, for all integers $\ell, n \geq 1$, iterating the above observation for n taking values $n, 2n, \ldots, 2^{\ell-1}n$ gives that

$$\mathbb{E}[V_{2^{\ell}n}] \ge \mathbb{E}[V_n] - K'' n^{-\delta} \sum_{j=0}^{\ell-1} 2^{-\delta j} \ge \mathbb{E}[V_n] - \frac{K''}{1 - 2^{-\delta}} n^{-\delta}.$$
 (7)

Taking $\ell \to \infty$, we conclude that $v_{\infty} \geq \mathbb{E}[V_n] - O(n^{-\delta})$. A similar reasoning exchanging Player 1 with Player 2 shows that $v_{\infty} \leq \mathbb{E}[V_n] + O(n^{-\delta})$ and concludes the proof of the lemma.

Finally, we are ready to prove Theorem 10.

Proof of Theorem 10. Fix an arbitrary $\varepsilon > 0, z \in Z$ and let $V_n = V_n(z)$. By Lemma 17, there is a constant K > 0 such that $|v_{\infty} - \mathbb{E}[V_n]| \leq Kn^{-\delta}$ for all $n \geq 1$. Combining this with the triangle inequality and Lemma 14 shows that, for every $t \geq 0$,

$$\mathbb{P}(|V_n - v_{\infty}| \ge t + Kn^{-\delta}) \le \mathbb{P}(|V_n - \mathbb{E}[V_n]| \ge t + Kn^{-\delta} - |\mathbb{E}[V_n] - v_{\infty}|)$$
$$\le \mathbb{P}(|V_n - \mathbb{E}[V_n]| \ge t) \le 2\exp\left(\frac{-t^2n^2}{2h(n)}\right),$$

which is the desired result.

5 Directed games on trees: proof of Theorem 11

The first lemma in this section bootstraps upon the conclusion of Lemma 14 (which still holds in this setting), thus deriving superexponential concentration for the value of the *n*-stage game. Below, log stands for the natural logarithm.

Lemma 18. Fix $\delta \in (0, 1/2)$, $n \ge 1$ and $t \ge n^{-\delta}$. For every even integer $k \in [2, n]$ such that

$$k \log d + 2 \log 2 \le t^2 (n - k),$$
 (8)

we have

$$\mathbb{P}(nV_n - (n-k)\mathbb{E}[V_{n-k}] \ge (n-k)t + k) \le \exp\left(-\frac{d^{k/2}}{6}\right),$$
$$\mathbb{P}(nV_n - (n-k)\mathbb{E}[V_{n-k}] \le -(n-k)t - k) \le \exp\left(-\frac{d^{k/2}}{6}\right).$$

The proof goes roughly as follows. Given an intermediate stage k = k(n) of the game such that (n - k(n)) grows to infinity, there typically exist only a few vertices such that the value of the (n - k)-game starting from them has a value that is very far from the expectation (we see these vertices as "bad"). Therefore, with foresight, both players can avoid such "bad" vertices, thus boosting the concentration of the *n*-value.

Proof of Lemma 18. First of all, since T is a transitive graph, for every fixed $n \ge 1$, the variables $(V_n(z))_z$ have the same distribution. For every even integer $k \in [n]$, denote

$$S_k \coloneqq \{z \in Z_k : V_{n-k}(z) - \mathbb{E}[V_{n-k}] \ge t\}.$$

In other words, S_k is the set of vertices z that could be reached from z_0 after k stages, for which the value of the (n-k)-stage game starting at z is greater than or equal to $\mathbb{E}[V_{n-k}] + t$.

Define the event $\mathcal{E}_k \coloneqq \{|S_k| \ge d^{k/2}\}$. We provide an upper bound for $\mathbb{P}(\mathcal{E}_k)$. Since the random variables $(V_{n-k}(z))_{z \in Z_k}$ are i.i.d., we have that $|S_k|$ follows a binomial distribution $\operatorname{Bin}(d^k, q)$ where $q \coloneqq \mathbb{P}(V_{n-k} \ge \mathbb{E}[V_{n-k}] + t)$. Consequently, by Lemma 14 (where the transient speed h is the identity function on N since, for all $z \in Z$ and $k \ge 2$, $Z_k(z)$ consists of all descendants of z at distance k, which form an antichain),

 $|S_k|$ is stochastically dominated by a binomial random variable $Bin(d^k, \tilde{q})$ where $\tilde{q} = \exp(-t^2(n-k)/2)$. In particular,

$$\mathbb{P}(\mathcal{E}_k) \leq \mathbb{P}\left(\operatorname{Bin}(d^k, \tilde{q}) \geq d^{k/2}\right).$$

The random variable $\operatorname{Bin}(d^k, \tilde{q})$ has mean $\mu \coloneqq d^k \tilde{q}$. We define

$$\xi \coloneqq \frac{d^{k/2}}{\mu} - 1 = \exp\left(\frac{t^2(n-k) - k\log(d)}{2}\right) - 1 \ge 1,$$

where the last inequality comes from (8). Since $d^{k/2} = (1+\xi)\mu$, we have that

$$\mathbb{P}(\operatorname{Bin}(d^k, \tilde{q}) \ge d^{k/2}) = \mathbb{P}(\operatorname{Bin}(d^k, \tilde{q}) \ge (1+\xi)\mu).$$

Therefore, since $\xi \ge 1$ (so $3\xi \ge 2 + \xi$), by Chernoff's bound,

$$\begin{split} \mathbb{P}\left(\mathrm{Bin}(d^{k},\tilde{q}) \geq d^{k/2}\right) &\leq \exp\left(-\frac{\xi^{2}\mu}{2+\xi}\right) \\ &\leq \exp\left(-\frac{\xi\mu}{3}\right) \\ &= \exp\left(-\frac{d^{k/2}}{3}\left(1-d^{k/2}\tilde{q}\right)\right) \,. \end{split}$$

Since $\xi = 1/(d^{k/2}\tilde{q}) - 1 \ge 1$, we have that $1 - d^{k/2}\tilde{q} \ge 1/2$, which finally yields

$$\mathbb{P}(\mathcal{E}_k) \le \exp\left(-\frac{d^{k/2}}{6}\right) \,. \tag{9}$$

At the same time, on the event $|S_k| < d^{k/2}$ (that is, $\overline{\mathcal{E}_k}$), Player 2 can ensure that the token avoids ending up in S_k after k stages. Indeed, at each of the $k/2 \in \mathbb{N}$ turns corresponding to decisions of Player 2, by the pigeonhole principle, Player 2 can always move the token to a vertex having at most a (1/d)-fraction of all remaining elements in S_k among its descendants. Since Player 2 has k/2 turns and $d^{-k/2}|S_k| < 1$, Player 2 can safely avoid the set S_k at stage k. Let us condition on the event $\overline{\mathcal{E}_k}$. Then, Player 2 can guarantee that the sum of the

Let us condition on the event \mathcal{E}_k . Then, Player 2 can guarantee that the sum of the payoffs over the last n-k stages is strictly smaller than $(n-k)(\mathbb{E}[V_{n-k}]+t)$. Moreover, the sum of the first k payoffs is at most k. Consequently, Player 2 can guarantee that, after n stages, the global mean payoff is strictly smaller than $k/n + (n-k)(\mathbb{E}[V_{n-k}]+t)/n$, in other words,

$$nV_n < (n-k)\mathbb{E}[V_{n-k}] + (n-k)t + k.$$
(10)

In particular, using (9) implies that

$$\mathbb{P}(nV_n - (n-k)\mathbb{E}[V_{n-k}] \ge (n-k)t + k) \le \mathbb{P}(|S_k| \ge d^{k/2}) = \mathbb{P}(\mathcal{E}_k) \le \exp\left(-\frac{d^{k/2}}{6}\right).$$

A similar reasoning for Player 1 (using the sets $\tilde{S}_k \coloneqq \{z \in Z_k : V_{n-k}(z) - \mathbb{E}[V_{n-k}] \le -t\}$ instead of S_k and replacing (10) with $nV_n > (n-k)\mathbb{E}[V_{n-k}] - (n-k)t$ yields

$$\mathbb{P}(nV_n - (n-k)\mathbb{E}[V_{n-k}] \le -(n-k)t) \le \exp\left(-\frac{d^{k/2}}{6}\right),$$

which implies the second statement. Note that the additional -k in it is introduced for reasons of symmetry only.

Next, we show that the expected value of the n-stage game converges rapidly as n grows to infinity.

Lemma 19. There exists $v_{\infty} \in \mathbb{R}$ such that, for every $\delta \in (0, 1/2)$, we have $|\mathbb{E}[V_n] - v_{\infty}| = O(n^{-\delta})$.

The arguments in the proof are very similar to the ones from the proof of Theorem 17 but the stronger concentration derived in Lemma 18 replaces the standard bounded difference estimate from Lemma 14.

Proof. Fix $\delta' \in (0, 1/2)$, $n \ge 1$, and $t \ge n^{-\delta'}$. Set $k = k(n) := 2\lfloor n^{1-2\delta'}/4 \log d \rfloor$. Then, $k \log d + 2 \log 2 \le t^2(n-k)$ for all large n. For every even integer $m \in [n/2, 2n]$ and large n, we have

$$\mathbb{P}\left(\min_{z\in Z_m} nV_n(z) \le (n-k)(\mathbb{E}[V_{n-k}]-t)-k\right)$$
$$\le \sum_{z\in Z_m} \mathbb{P}\left(nV_n(z) \le (n-k)(\mathbb{E}[V_{n-k}]-t)-k\right)$$
$$\le d^m \exp\left(-\frac{d^{k/2}}{6}\right)$$
$$\le \exp\left(2n\log d - \frac{d^{\lfloor n^{1-2\delta'}/4\log d\rfloor}}{6}\right),$$

where the first inequality comes from a union bound and the second inequality comes from Theorem 18. Fix $\delta \in (0, \delta')$ and define, for all $n \ge 1$,

$$\varepsilon_n \coloneqq n^{-\delta}$$
 and $\psi(n) \coloneqq \exp\left(2n\log d - d^{\lfloor n^{1-2\delta'}/4\log d\rfloor}/6\right)$.

For large n and every even integer $m \in [n/2, 2n]$, we have

$$\mathbb{E}\left[\min_{z\in\mathbb{Z}_m}V_n(z)\right] \ge \left(\frac{n-k}{n}(\mathbb{E}[V_{n-k}] - \varepsilon_n) - \frac{k}{n}\right)$$
(11)

$$\mathbb{P}\left(\min_{z\in\mathbb{Z}_{m}}nV_{n}(z) > (n-k)(\mathbb{E}[V_{n-k}]-\varepsilon_{n})-k\right) \\
\geq \left(\frac{n-k}{n}(\mathbb{E}[V_{n-k}]-\varepsilon_{n})-\frac{k}{n}\right)(1-\psi(n)) \\
\geq \left(\mathbb{E}[V_{n-k}]-\frac{k}{n}(1+\mathbb{E}[V_{n-k}])-\varepsilon_{n}\right)(1-\psi(n)) \\
\geq \left(\mathbb{E}[V_{n}]-\frac{3k}{n}-\varepsilon_{n}\right)(1-\psi(n)) \geq \mathbb{E}[V_{n}]-(\psi(n)+2\varepsilon_{n}), \quad (12)$$

where in the fourth inequality we used that $\mathbb{E}[V_n] \leq \mathbb{E}[V_{n-k}] + k/n$ by Lemma 15 and $1 + \mathbb{E}[V_{n-k}] \leq 2$, and the last inequality is valid for large *n* because $k/n = o(\varepsilon_n)$.

Consider integers $n \ge 1$ and even $m \in [n/2, 2n]$. In the (n+m)-stage game, Player 1 can play according to an optimal strategy for the *m*-stage game starting at z_0 , and then play according to an optimal strategy for the *n*-stage game starting from the state z reached after *m* stages. This guarantees that $(m+n)V_{m+n} \ge mV_m + \min_{z \in Z_m} nV_n(z)$. Taking expectations on both sides and using (12) yields

$$(m+n)\mathbb{E}[V_{m+n}] \ge m\mathbb{E}[V_m] + n\mathbb{E}\left[\min_{z\in Z_m} V_n(z)\right]$$
$$\ge m\mathbb{E}[V_m] + n\mathbb{E}[V_n] - n(\psi(n) + 2\varepsilon_n).$$

We find a similar inequality for odd $m \in [n/2, 2n]$. In this case, m + 1 is even and also in [n/2, 2n]. Then, the previous inequality applied to m + 1 and n yields

$$(m+n+1)\mathbb{E}[V_{m+n+1}] \ge (m+1)\mathbb{E}[V_{m+1}] + n\mathbb{E}[V_n] - n(\psi(n) + 2\varepsilon_n).$$
(13)

However,

$$(m+n)\mathbb{E}[V_{m+n}] \ge (m+n+1)\mathbb{E}[V_{m+n+1}] - 1$$
 and $(m+1)\mathbb{E}[V_{m+1}] \ge m\mathbb{E}[V_m]$,

which combined with (13) gives

$$(m+n)\mathbb{E}[V_{m+n}] \ge m\mathbb{E}[V_m] + n\mathbb{E}[V_n] - n(\psi(n) + 2\varepsilon_n) - 1.$$

To sum things up, for large n and $m \in [n/2, 2n]$,

$$(m+n)\mathbb{E}[V_{m+n}] \ge m\mathbb{E}[V_m] + n\mathbb{E}[V_n] - n(\psi(n) + 2\varepsilon_n) - 1.$$
(14)

Recall that there is a constant K' > 0 such that, for all $n \ge 1$, $n(\psi(n) + 2\varepsilon_n) + 1 \le K' n^{1-\delta}$. We define $\phi(n) \coloneqq K' n^{1-\delta}$ and deduce from (14) that

$$(m+n)\mathbb{E}[V_{m+n}] \ge m\mathbb{E}[V_m] + n\mathbb{E}[V_n] - \phi(n+m).$$

Moreover, ϕ is increasing and verifies $\sum_{n\geq 1} \phi(n)/n^2 < \infty$. Consequently, Lemma 13 applied to the function $f: n \in \mathbb{N} \mapsto -n\mathbb{E}[V_n]$ implies that that $\mathbb{E}[V_n]$ converges to a limit $v_{\infty} \in \mathbb{R} \cup \{\infty\}$ as $n \to \infty$. Note that $v_{\infty} \in [0, 1]$ since $V_n \in [0, 1]$ for all $n \geq 1$.

Finally, using (14) with m = n and a telescopic summation shows that the inequality (7) still holds. In particular, we conclude that $v_{\infty} \geq \mathbb{E}[V_n] - O(n^{-\delta})$. A similar reasoning replacing Player 1 with Player 2 shows that $v_{\infty} \leq \mathbb{E}[V_n] + O(n^{-\delta})$ and concludes the proof of the lemma.

We are now ready to prove Theorem 11.

Proof of Theorem 11. Fix $t \ge n^{-\delta}$ and let K' be a constant such that $|\mathbb{E}[V_n] - v_{\infty}| \le K'n^{-\delta}$ for all large n. Using that, for all n and $k \le n$, we have $|nV_n - (n-k)V_{n-k}| \le k$, and fixing $k = 2\lceil \frac{t^2n}{4\log d} \rceil$ (which satisfies (8)), we get

$$\begin{split} \mathbb{P}(|V_n - v_{\infty}| \geq t + 2t^2 + K'n^{-\delta}) \\ &\leq \mathbb{P}\bigg(\left|V_n - \frac{n-k}{n}\mathbb{E}[V_{n-k}]\right| \geq t + 2t^2 + K'n^{-\delta} \\ &- \left|\frac{n-k}{n}\mathbb{E}[V_{n-k}] - \mathbb{E}[V_n]\right| - |\mathbb{E}[V_n] - v_{\infty}|\bigg) \\ &\leq \mathbb{P}\left(\left|V_n - \frac{n-k}{n}\mathbb{E}[V_{n-k}]\right| \geq t + t^2\right) \\ &\leq \mathbb{P}(|nV_n - (n-k)\mathbb{E}[V_{n-k}]| \geq (n-k)t + k) \\ &\leq \exp\left(-\frac{d^{\lfloor k/2 \rfloor}}{6}\right) \leq \exp\left(-\frac{d^{t^2n/(4\log d)}}{6}\right) = \exp\left(-\frac{1}{6}\exp\left(\frac{t^2n}{4}\right)\right), \end{split}$$

where the first inequality comes from the triangle inequality, the second inequality comes from the definition of K' and the fact that $|nV_n - (n-k)V_{n-k}| \le k \le nt^2$, and the third inequality once again uses the fact that $k \le nt^2$.

Finally, choosing $K \ge K'$ sufficiently large ensures that, first, the upper bound shown above holds for all $n \ge 1$ (and not only for large n), and second, the upper bound holds for all $t \ge 0$, which finishes the proof.

6 Oriented games: proof of Proposition 6

We present a simple and self-contained proof of Proposition 6.

Proof of Proposition 6. First, by density of the rational vectors in \mathbb{R}^d and rescaling, we may assume that $u \in \mathbb{Z}^d$ is such that the greatest common divisor of its coordinates is 1. We provide an upper bound on h(n). For every integer $i \ge 1$ and initial state $z_0 = z$, define $Z_{2i}(z) \coloneqq \{w \in Z : w \cdot u = z \cdot u + i\}$, $Z_{2i+1}(z) \coloneqq \{w \in Z : w \cdot u = z \cdot u - i\}$, and $Z_0(z) \coloneqq \{z\}$ and $Z_1(z) \coloneqq \{w \in Z \setminus \{z\} : w \cdot u = z \cdot u\}$. Then, for all $z \in Z$, $(Z_i(z))_{i\ge 1}$ form a partition of Z and each of them could be visited at most once, i.e., for all $n \ge 1$, the sets (Z_1, Z_2, \ldots, Z_n) forms an antichain. Set $r \coloneqq \max_{(u,v) \in E(\Gamma)} \|v - u\|_2$. After n

steps of the game, the position z_n satisfies $||z_n - z||_2 \le nr$, and by the Cauchy-Schwarz inequality,

$$|(z_n - z) \cdot u| \le ||z_n - z||_2 \cdot ||u||_2 \le \lceil nr \cdot ||u||_2 \rceil =: N = N(n).$$

In particular, $Z^{(n)}(z)$ is contained in the ball with radius N around z, which itself is contained in $Z_1 \cup Z_2 \cup \ldots \cup Z_{2N+1}(z)$, so the transient speed of the process satisfies $h(n) \leq 2N(n) + 1$ for all $n \geq 1$.

Now, fix $\delta \in (0, 1/2)$ and $z_0 = z \in Z$. We show that the game is δ -transient. Set $\varepsilon_n \coloneqq n^{-\delta}$. Then,

$$\psi(n,\varepsilon_n) = \exp\left(-\frac{\varepsilon_n^2 n^2}{2h(n)}\right) \max_{z\in Z} |Z^{(2n)}(z)|$$

$$\leq \exp\left(-\frac{\varepsilon_n^2 n}{6r \cdot ||u||_2}\right) (2nr \cdot ||u||_2 + 1)^d$$

$$= \exp\left(-\frac{n^{1-2\delta}}{6r \cdot ||u||_2}\right) (2nr \cdot ||u||_2 + 1)^d = O(n^{-\delta})$$

Hence, for all $\delta \in (0, 1/2)$, $\varepsilon_n + \psi(n, \varepsilon_n) = O(n^{-\delta})$, and therefore, the game is δ -transient.

Declarations

Funding

This work was supported by the French Agence Nationale de la Recherche (ANR) under references ANR-21-CE40-0020 (CONVERGENCE project), ANR-20-CE40-0002 (GrHyDy), and under grant ANR-17-EURE-0010 (Programme Investissements d'Avenir), by Fondecyt grant 1220174, by ANID Chile grant ACT210005, and by the ERC CoG 863818 (ForM-SMArt) grant. This collaboration was mainly conducted during a 1-year visit of Bruno Ziliotto to the Center for Mathematical Modeling (CMM) at University of Chile in 2023, under the IRL program of CNRS. This work was supported by Fondation CFM pour la Recherche.

Conflict of interest/Competing interests

The authors have no competing interests to declare that are relevant to the content of this article.

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Data availability

Not applicable.

Code availability

Not applicable.

References

- Ehrenfeucht, A., Mycielski, J.: Positional strategies for mean payoff games. International Journal of Game Theory 8(2), 109–113 (1979) https://doi.org/10.1007/ BF01768705
- Shapley, L.S.: Stochastic Games. Proceedings of the National Academy of Sciences 39(10), 1095–1100 (1953) https://doi.org/10.1073/pnas.39.10.1095
- Bewley, T., Kohlberg, E.: The Asymptotic Theory of Stochastic Games. Mathematics of Operations Research 1(3), 197–208 (1976) https://doi.org/10.1287/moor.1.3.197 3689563
- [4] Laraki, R., Sorin, S.: Advances in zero-sum dynamic games. In: Handbook of Game Theory with Economic Applications vol. 4, pp. 27–93. Elsevier, ??? (2015)
- [5] Levy, Y., Solan, E.: Stochastic games. Complex Social and Behavioral Systems: Game Theory and Agent-Based Models, 229–250 (2020)
- [6] Solan, E.: A Course in Stochastic Game Theory vol. 103. Cambridge University Press, ??? (2022)
- [7] Solan, E., Ziliotto, B.: Stochastic games with signals. Advances in Dynamic and Evolutionary Games: Theory, Applications, and Numerical Methods, 77–94 (2016)
- [8] Garrec, T.: Communicating zero-sum product stochastic games. Journal of Mathematical Analysis and Applications 477(1), 60–84 (2019)
- [9] Laraki, R., Renault, J.: Acyclic gambling games. Mathematics of Operations Research 45(4), 1237–1257 (2020)
- [10] Ziliotto, B.: Mertens conjectures in absorbing games with incomplete information. The Annals of Applied Probability 34(2) (2024) https://doi.org/10.1214/ 23-AAP2011
- [11] Ziliotto, B.: Zero-sum repeated games: counterexamples to the existence of the asymptotic value and the conjecture maxmin $= \lim v(n)$. The Annals of Probability 44(2), 1107–1133 (2016)

- [12] Garnier, G., Ziliotto, B.: Percolation games. Mathematics of Operations Research 48(4), 2156–2166 (2023)
- [13] Alon, N., Rudov, K., Yariv, L.: Dominance Solvability in Random Games (2021)
- [14] Amiet, B., Collevecchio, A., Scarsini, M., Zhong, Z.: Pure Nash equilibria and best-response dynamics in random games. Mathematics of Operations Research 46(4), 1552–1572 (2021)
- [15] Flesch, J., Predtetchinski, A., Suomala, V.: Random perfect information games. Mathematics of Operations Research 48(2), 708–727 (2023)
- [16] Heinrich, T., Jang, Y., Mungo, L., Pangallo, M., Scott, A., Tarbush, B., Wiese, S.: Best-response dynamics, playing sequences, and convergence to equilibrium in random games. International Journal of Game Theory 52(3), 703–735 (2023) https://doi.org/10.1007/s00182-023-00837-4
- [17] Holroyd, A.E., Marcovici, I., Martin, J.B.: Percolation games, probabilistic cellular automata, and the hard-core model. Probability Theory and Related Fields 174, 1187–1217 (2019)
- [18] Bhasin, D., Karmakar, S., Podder, M., Roy, S.: On a class of PCA with size-3 neighborhood and their applications in percolation games. Electronic Journal of Probability 28, 1–60 (2023)
- [19] Ziliotto, B.: Stochastic homogenization of nonconvex hamilton-jacobi equations: A counterexample. Communications on Pure and Applied Mathematics (2017)
- [20] Davini, A., Saona, R., Ziliotto, B.: Stochastic homogenization of HJ equations: a differential game approach (2024)
- [21] Kuhn, H.W.: Extensive games. Proceedings of the National Academy of Sciences of the United States of America 36(10), 570–576 (1950) 88394
- [22] Balakrishnan, K., Lekha, D.S., Changat, M., Anand, B.S., Narasimha-Shenoi, P.G.: Generalized Vertex Transitivity in Graphs (2018)
- [23] Mirsky, L.: A Dual of Dilworth's Decomposition Theorem. The American Mathematical Monthly 78(8), 876–877 (1971) https://doi.org/10.1080/00029890.1971. 11992886
- [24] Janson, S., Łuczak, T., Ruciński, A.: Random Graphs. John Wiley & Sons, Inc., ??? (2000)
- [25] Bruijn, N.G., Erdős, P.: Some linear and some quadratic recursion formulas II. Indagationes Mathematicae (Proceedings) 55, 152–163 (1952) https://doi.org/ 10.1016/S1385-7258(52)50021-0