8-2021

Bidding mechanisms in graph games

Guy AVNI

Thomas A. HENZINGER

Dorde ZIKELIC
Singapore Management University, dzikelic@smu.edu.sg

Citation

AVNI, Guy; HENZINGER, Thomas A.; and ZIKELIC, Dorde. Bidding mechanisms in graph games. (2021). *Journal of Computer and System Sciences*. 119, 133-144.
Available at: https://ink.library.smu.edu.sg/sis_research/9060

Journal of Computer and System Sciences 119 (2021) 133–144

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Bidding mechanisms in graph games

Guy Avni [a,*], Thomas A. Henzinger [b], Đorđe Žikelić [b]

[a] University of Haifa, Israel
[b] IST Austria, Austria

ARTICLE INFO

Article history:
Received 12 October 2020
Received in revised form 16 January 2021
Accepted 25 February 2021
Available online 3 March 2021

Keywords:
Graph games
Bidding games
Richman bidding
Poorman bidding
Mean-payoff games
Parity games
Stochastic games
Random-turn games

ABSTRACT

A graph game proceeds as follows: two players move a token through a graph to produce a finite or infinite path, which determines the payoff of the game. We study *bidding games* in which in each turn, an auction determines which player moves the token. Bidding games were largely studied in combination with two variants of first-price auctions called "Richman" and "poorman" bidding. We study *taxman* bidding, which span the spectrum between the two. The game is parameterized by a constant $\tau \in [0, 1]$: portion τ of the winning bid is paid to the other player, and portion $1 - \tau$ to the bank. While finite-duration (reachability) taxman games have been studied before, we present, for the first time, results on *infinite-duration* taxman games: we unify, generalize, and simplify previous equivalences between bidding games and a class of stochastic games called *random-turn games*.

1. Introduction

Bidding games on graphs. Two-player infinite-duration games on graphs are a central class of games in formal verification [1], where they are used, for example, to solve synthesis [2], and they have deep connections to foundations of logic [3]. A graph game proceeds by placing a token on a vertex in the graph, which the players move throughout the graph to produce an infinite path ("play") π. The game is zero-sum and π determines the winner or payoff. Graph games can be classified according to the players' objectives. For example, the simplest objective is *reachability*, where Player 1 wins iff an infinite path visits a designated target vertex. Another classification of graph games is the *mode of moving* the token. The most studied mode of moving is *turn based*, where the players alternate turns in moving the token.

In *bidding games*, in each turn, an "auction" is held between the two players in order to determine which player moves the token. Bidding games were first introduced in [4,5], where the following variants of *first-price* auctions were considered. Each player has a budget, and before each turn, the players submit bids simultaneously, where a bid is legal if it does not exceed the available budget, and the higher bidder moves the token. The variants differ in the recipient of the bid:

- In *Richman* bidding (named after David Richman), the higher bidder pays the lower bidder.
- In *poorman* bidding, the higher bidder pays the bank, thus the money is lost.
- *Taxman* bidding spans the spectrum between Richman and poorman. It is parameterized by $\tau \in [0, 1]$: the winner of a bidding pays portion τ of his bid to the other player and portion $1 - \tau$ to the bank.

* Corresponding author.
E-mail address: gavni@cs.haifa.ac.il (G. Avni).

https://doi.org/10.1016/j.jcss.2021.02.008

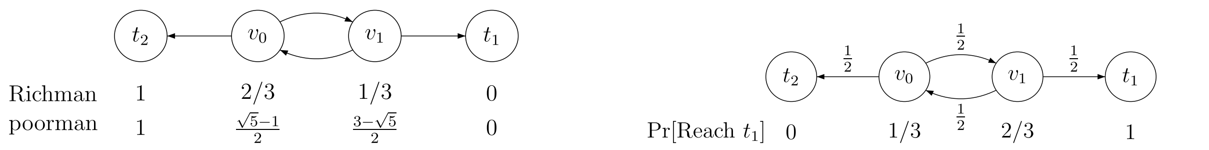

Fig. 1. Left: A reachability bidding game $\mathcal{G}$ in which the target of Player i, for $i \in \{1, 2\}$, is t_i. The threshold ratios under Richman and poorman bidding appear below the vertices. For example, under Richman bidding it is convenient to normalize the sum of budgets to 1. Then, Player 1 wins as follows from v_0 with a budget of $\frac{2}{3} + \epsilon$, for $\epsilon > 0$. He bids $\frac{1}{3}$ at v_0, thus necessarily overbidding Player 2. He bids his whole budget at v_1, thus either winning the game, or returning with a budget of $\frac{2}{3} + 2\epsilon$ to v_0. This strategy ensures that Player 1's budget eventually exceeds $\frac{3}{4}$, which allows him to force two consecutive bidding wins and forcing the game from v_0 to t_1. **Right:** The (simplified) unbiased random-turn game $\texttt{RT}(\mathcal{G}, 0.5)$ is a Markov chain. The value of a vertex is the probability of reaching t_1. Note that under Richman bidding, for every vertex v, we have $\texttt{Th}(\mathcal{G}, v) = 1 - val\big(\texttt{RT}(\mathcal{G}, 0.5), v\big)$. Moreover, under poorman bidding, ratios are irrational thus such an equivalence is unlikely to exist.

Note that taxman bidding with $\tau = 1$ coincides with Richman bidding and taxman bidding with $\tau = 0$ coincides with poorman bidding. A central quantity in bidding games is the *initial ratio* of the players budgets. Formally, assuming that, for $i \in \{1, 2\}$, Player i's initial budget is B_i, we say that Player 1's initial ratio is $B_1/(B_1 + B_2)$.

Reachability bidding games. Reachability bidding games were studied in [4], where the central question that was studied regards the existence of a necessary and sufficient initial ratio to guarantee winning the game. Formally, the *threshold ratio* in a vertex v, denoted $\texttt{Th}(v)$, is such that if Player 1's initial ratio exceeds $\texttt{Th}(v)$, he can guarantee winning the game, and if his initial ratio is less than $\texttt{Th}(v)$, Player 2 can guarantee winning the game.[1] Existence of threshold ratios in reachability games for all three bidding mechanisms was shown in [4].

An interesting equivalence was observed in [4] between reachability Richman-bidding games (and only them) with a class of games called *random-turn games* [6] (see Fig. 1), which are a subclass of stochastic games [7]. A random-turn game is parameterized by $p \in [0, 1]$. In each turn, rather than bidding, the player who moves is chosen by throwing a (possibly) biased coin: with probability p, Player 1 chooses how to move the token, and Player 2 chooses with probability $1 - p$. Consider a reachability Richman-bidding game $\mathcal{G}$. We construct a "uniform" random-turn game on top of $\mathcal{G}$, denoted $\texttt{RT}(\mathcal{G}, 0.5)$, in which we throw an unbiased coin in each turn. The objective of Player 1 remains reaching his target vertex. It is well known that each vertex in $\texttt{RT}(\mathcal{G}, 0.5)$ has a *value*, which is, informally, the probability of reaching the target when both players play optimally, and which we denote by $val\big(\texttt{RT}(\mathcal{G}, 0.5), v\big)$. The equivalence relates the threshold budget in a Richman bidding game with the value in a random-turn game. Specifically, for every vertex v in the Richman game $\mathcal{G}$, the threshold ratio in v equals $1 - val\big(\texttt{RT}(\mathcal{G}), v\big)$. We note that such a connection is not known and is unlikely to exist in reachability games with neither poorman nor taxman bidding since threshold ratios can be irrational.

Infinite-duration bidding games. Infinite-duration bidding games have been recently studied with Richman [8] and poorman [9] bidding. For qualitative objectives, namely games in which one player wins and the other player loses, both bidding rules have similar properties. By reducing general qualitative games to reachability games, it is shown that threshold ratios exist for both types of bidding rules. We show a similar result for qualitative games with taxman bidding.

Things get interesting in *mean-payoff* games, which are quantitative games: an infinite play has a *payoff*, which is Player 1's reward and Player 2's cost (see an example of a mean-payoff game in Fig. 2). We thus call the players in a mean-payoff game Max and Min, respectively. Our most interesting results regard games played on strongly-connected graphs. Consider a strongly-connected mean-payoff bidding game $\mathcal{G}$. For $p \in [0, 1]$, the random-turn game $\texttt{RT}(\mathcal{G}, p)$ is a mean-payoff stochastic game. Its *value*, denoted $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, p)\big)$, is again a well-known concept and equals the expected payoff under optimal play of the two players.

Mean-payoff Richman bidding games are equivalent to uniform random-turn games [8]. More formally, in a strongly-connected game $\mathcal{G}$, irrespective of the initial ratio and for every $\epsilon > 0$, Max can guarantee a payoff of $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, 0.5)\big) - \epsilon$ and Min can guarantee a payoff of at most $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, 0.5)\big) + \epsilon$. With poorman bidding, contrary to Richman bidding, the optimal payoff depends on the initial ratio. That is, with a higher initial ratio, Max can guarantee a better payoff. Moreover, a strongly-connected mean-payoff game $\mathcal{G}$ with initial ratio $r \in (0, 1)$ for Max is equivalent to the random-turn game that uses a coin with a bias of r towards Max. That is, the optimal payoff Max can guarantee in $\mathcal{G}$ with an initial ratio r is roughly $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, r)\big)$. We find this equivalence particularly surprising given that no equivalence is known for reachability poorman bidding games.

Mean-payoff taxman-bidding games. Recall that taxman bidding spans the spectrum between Richman and poorman bidding. The differences between Richman and poorman bidding, raise the question of identifying the properties of mean-payoff taxman bidding games. Our main contribution is a complete solution to this question.

[1] When the initial ratio is exactly $\texttt{Th}(v)$, the winner depends on the mechanism with which ties are broken. Our results do not depend on a specific tie-breaking mechanism.

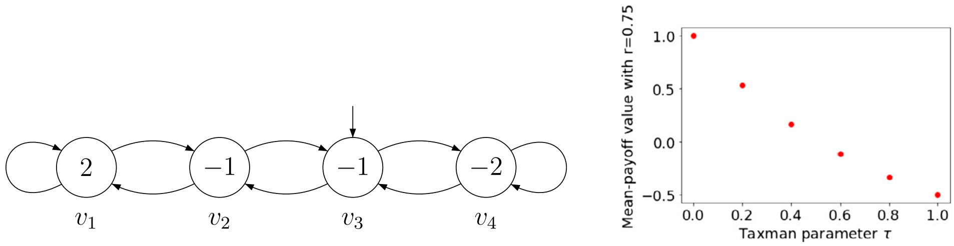

Fig. 2. On the left, a mean-payoff game $\mathcal{G}$. On the right, the mean-payoff value of $\mathcal{G}$, where the initial ratio is fixed to 0.75 and the taxman parameter τ varies. The value of $\mathcal{G}$ with Richman bidding is -0.5, with poorman bidding, it is 1, and, for example, with $\tau = 0.2$, it is 0.533.

Theorem 1 *(Informal). A strongly-connected mean-payoff game $\mathcal{G}$, an initial ratio $r \in (0,1)$ for Max, and a taxman parameter $\tau \in [0,1]$ is equivalent to the random-turn game $RT(\mathcal{G}, \frac{r+\tau\cdot(1-r)}{1+\tau})$.*

As a sanity check, note that for $\tau = 1$, we have $\frac{r+\tau\cdot(1-r)}{1+\tau} = 0.5$, which agrees with the result on Richman bidding, and for $\tau = 0$, we have $\frac{r+\tau\cdot(1-r)}{1+\tau} = r$, which agrees with the result on poorman bidding. The solution to general mean-payoff taxman games follows from the solution to strongly-connected games. In Fig. 2, we depict some mean-payoff values for a fixed initial ratio and varying taxman parameter. Previous results only give the two endpoints in the plot, and the mid points in the plot are obtained using the results in this paper.

The main technical challenge is constructing an optimal deterministic strategy for Max. The construction of Max's strategy involves two components. First, we assign an "importance" to each vertex v, which we call *strength* and denote $\mathrm{St}(v)$. Intuitively, if $\mathrm{St}(v) > \mathrm{St}(u)$, then it is more important for Max to move in v than in u. Second, when the game reaches a vertex v, Max's bid is a careful normalization of $\mathrm{St}(v)$ so that changes in Max's ratio are matched with the accumulated weights in the game. Finding the right normalization is intricate and it constitutes the main technical contribution of this paper. Previous such normalizations were constructed for Richman and poorman mean-payoff games [8,9]. The construction for Richman bidding is much more complicated than the one we present here. The construction for poorman bidding is ad-hoc and does not generalize. Our construction for taxman bidding thus unifies these constructions and simplifies them. It uses techniques that can generalize beyond taxman bidding. Moreover, the strategies that we construct give guarantees for finite prefixes: our strategies maintain a fixed ratio between the number bidding wins and loses up to an additive constant that depends on the initial ratio and arena. Finally, we study, for the first time, complexity problems for taxman games.

2. Preliminaries

A graph game is played on a directed graph $G = \langle V, E\rangle$, where V is a finite set of vertices and $E \subseteq V \times V$ is a set of edges. The *neighbors* of a vertex $v \in V$, denoted $N(v)$, is the set of vertices $\{u \in V : \langle v, u\rangle \in E\}$. A *path* in G is a finite or infinite sequence of vertices $v_1, v_2, \ldots$ such that for every $i \geq 1$, we have $\langle v_i, v_{i+1}\rangle \in E$.

Bidding games. Each Player i has a budget $B_i \in \mathbb{R}^{\geq 0}$. In each turn a bidding determines which player moves the token. Both players simultaneously submit bids, where a bid b_i for Player i is legal if $b_i \leq B_i$. The player who bids higher wins the bidding, where we assume some mechanism to break ties, e.g., always giving Player 1 the advantage, and our results are not affected by the specific tie-breaking mechanism at use. The winner moves the token and pays his bid, where we consider three bidding mechanisms that differ in where the winning bid is paid. Suppose Player 1 wins a bidding with his bid of b.

- In **Richman** bidding, the winner pays to the loser, thus the new budgets are $B_1 - b$ and $B_2 + b$.
- In **poorman** bidding, the winner pays to the bank, thus the new budgets are $B_1 - b$ and B_2.
- In **taxman** bidding with parameter $\tau \in [0,1]$, the winner pays portion τ to the other player and $(1-\tau)$ to the bank, thus the new budgets are $B_1 - b$ and $B_2 + \tau \cdot b$.

A central quantity in bidding games is the *ratio* of a player's budget from the total budget.

Definition 2 *(**Ratio**).* Suppose the budget of Player i is B_i, for $i \in \{1,2\}$, at some point in the game. Then, Player i's *ratio* is $B_i/(B_1 + B_2)$. The *initial ratio* refers to the ratio of the initial budgets, namely the budgets before the game begins. We restrict attention to games in which both players start with positive initial budgets, thus the initial ratio is in $(0,1)$.

Strategies and plays. A *strategy* is a recipe for how to play a game. It is a function that, given a finite *history* of the game, prescribes to a player which *action* to take, where we define these two notions below. For example, in turn-based games, a strategy takes as input, the sequence of vertices that were visited so far, and it outputs the next vertex to move to. In bidding games, histories and strategies are more involved as they maintain the information about the bids and winners of the bids. Formally, a history in a bidding game is $\pi = \langle v_1, b_1, i_1\rangle, \ldots, \langle v_k, b_k, i_k\rangle, v_{k+1} \in (V \times \mathbb{R} \times \{1,2\})^* \cdot V$, where for

$1 \leq j \leq k+1$, the token is placed on vertex v_j at round j, for $1 \leq j \leq k$, the winning bid is b_j and the winner is Player i_j. Consider a finite history π. For $i \in \{1,2\}$, let $W_i(\pi) \subseteq \{1,\ldots,k\}$ denote the indices in which Player i is the winner of the bidding in π. Let B_i^I be the initial budget of Player i. Player i's budget following π, denoted $B_i(\pi)$, depends on the bidding mechanism. For example, in Richman bidding, $B_1(\pi) = B_i^I - \sum_{j \in W_1(\pi)} b_j + \sum_{j \in W_2(\pi)} b_j$, B_2 is defined dually, and the definition is similar for taxman and poorman bidding. Given a history π that ends in v, a strategy for Player i prescribes an action $\langle b, v \rangle$, where $b \leq B_i(\pi)$ is a bid that does not exceed the available budget and v is a vertex to move to upon winning, where we require that v is a neighbor of v_{k+1}. An initial vertex, initial budgets, and two strategies for the players determine a unique infinite *play* π for the game. The vertices that π visits form an infinite path $path(\pi)$.

Objectives. An objective O is a set of infinite paths. Player 1 wins an infinite play π iff $path(\pi) \in O$. We call a strategy f *winning* for Player 1 w.r.t. an objective O if for every strategy g of Player 2 the play that f and g determine is winning for Player 1. Winning strategies for Player 2 are defined dually. We consider the following qualitative objectives:

1. In *reachability games*, Player 1 has a target vertex t and an infinite play is winning iff it visits t.
2. In *parity games*, each vertex is labeled with an index in $\{1,\ldots,d\}$. An infinite path is winning for Player 1 iff the parity of the maximal index that is visited infinitely often is odd.
3. *Mean-payoff games* are played on weighted directed graphs, with weights given by a function $w : V \to \mathbb{Q}$. Consider an infinite path $\eta = v_1, v_2, \dots \in V^\omega$. For $n \in \mathbb{N}$, the prefix of length n of η is η^n, and we define its *energy* to be $E(\eta^n) = \sum_{i=1}^n w(v_i)$. The *payoff* of η is $\mathtt{MP}(\eta) = \liminf_{n\to\infty} E(\eta^n)/n$. Player 1 wins η iff $\mathtt{MP}(\eta) \geq 0$.

Mean-payoff games are quantitative games. We think of the payoff as Player 1's reward and Player 2's cost, thus in mean-payoff games, we refer to Player 1 as Max and to Player 2 as Min.

Threshold ratios. The first question that arises in the context of bidding games asks what is the necessary and sufficient initial ratio to guarantee an objective.

Definition 3 *(**Threshold ratios**).* Consider a bidding game $\mathcal{G}$, a vertex v, an initial ratio r, and an objective O for Player 1. The threshold ratio in v, denoted $\mathtt{Th}(v)$, is a ratio in $[0,1]$ such that if $r > \mathtt{Th}(v)$, then Player 1 has a winning strategy that guarantees that O is satisfied, and if $r < \mathtt{Th}(v)$, then Player 2 has a winning strategy that violates O.

Random-turn games. A *stochastic game* [7] is a graph game in which the vertices are partitioned between two players and a *nature* player. As in turn-based games, whenever the game reaches a vertex that is controlled by Player i, for $i = 1,2$, he choses how the game proceeds, and whenever the game reaches a vertex v that is controlled by nature, the next vertex is chosen according to a probability distribution that depends only on v.

Consider a bidding game $\mathcal{G}$ that is played on a graph $\langle V, E \rangle$. The *random-turn game* with bias $p \in [0,1]$ that is associated with $\mathcal{G}$ is a stochastic game that intuitively simulates the following process. In each turn we throw a biased coin that turns heads with probability p and tails with probability $1-p$. If the coin turns heads, then Player 1 moves the token, and otherwise Player 2 moves the token. Formally, we define $\mathtt{RT}(\mathcal{G}, p) = \langle V_1, V_2, V_N, E, \Pr \rangle$, where each vertex in V is split into three vertices, each controlled by a different player, thus for $\alpha \in \{1,2,N\}$, we have $V_\alpha = \{v_\alpha : v \in V\}$, nature vertices simulate the fact that Player 1 chooses the next move with probability p, thus $\Pr[v_N, v_1] = p = 1 - \Pr[v_N, v_2]$, and reaching a vertex that is controlled by one of the two players means that he chooses the next move, thus $E = \{\langle v_\alpha, u_N \rangle : \langle v, u \rangle \in E \text{ and } \alpha \in \{1,2\}\}$. When $\mathcal{G}$ is a mean-payoff game, the vertices are weighted and we define the weights of v_1, v_2, and v_N to be equal to the weight of v.

The following definitions are standard, and we refer the reader to [10] for more details. A strategy in a stochastic game is similar to a turn-based game; namely, given the history of vertices visited so far, the strategy chooses the next vertex. Fixing two such strategies f and g for both players gives rise to a distribution $D(f,g)$ on infinite paths. Intuitively, Player 1's goal is to maximize the probability that his objective is met. An *optimal* strategy for Player 1 guarantees that the objective is met with probability at least c and, intuitively, he cannot do better, thus Player 2 has a strategy that guarantees that the objective is violated with probability at least $(1-c)$. It is well known that optimal positional strategies exist for the objectives that we consider.

Definition 4 *(**Values in stochastic games**).* Consider a bidding game $\mathcal{G}$, let $p \in [0,1]$, and consider two optimal strategies f and g for the two players in $\mathtt{RT}(\mathcal{G}, p)$. When $\mathcal{G}$ is a qualitative game with objective O, the *value* of $\mathtt{RT}(\mathcal{G}, p)$, denoted $val\big(\mathtt{RT}(\mathcal{G}, p)\big)$, is $\Pr_{\eta \sim D(f,g)} \Pr[\eta \in O]$. When $\mathcal{G}$ is a mean-payoff game, the *mean-payoff value* of $\mathtt{RT}(\mathcal{G}, p)$, denoted $\mathtt{MP}\big(\mathtt{RT}(\mathcal{G}, p)\big)$, is $\mathbb{E}_{\eta \in D(f,g)} \mathtt{MP}(\eta)$.

3. Qualitative taxman games

In this section we survey known results on reachability bidding games and then study parity taxman games. The solution is based on a reduction from parity to reachability taxman games: we show in Lemma 6 that in each bottom strongly-connected component (BSCC), one of the players wins with any positive initial budget, and then construct a reachability game in which each player tries to force the game into a winning BSCC.

In [5,4], reachability bidding games were studied with a slightly different definition, which we call *double-reachability*: both players have a target, where we denote by t_i the target of Player i, for $i \in \{1,2\}$, all vertices have a path to both targets, and the game ends once one of the targets is reached. They show the following results.

Theorem 5. *[5,4] Consider a double-reachability bidding game $\mathcal{G}$ and a vertex v. The threshold ratio exists in v with Richman, poorman, and taxman bidding. Moreover, threshold ratios have the following properties. For the target vertex t_1 of Player 1, we have $\texttt{Th}(t_1) = 0$, and for the target t_2 of Player 2, we have $\texttt{Th}(v) = 1$. Consider some other vertex v and denote $v^+, v^- \in N(v)$ the vertices with the minimal and maximal thresholds in the neighborhood of v, thus for every $u \in N(v)$, we have $\texttt{Th}(v^-) \leq \texttt{Th}(u) \leq \texttt{Th}(v^+)$.*

- *In Richman bidding, we have $\texttt{Th}(v) = \frac{1}{2}\big(\texttt{Th}(v^+) + \texttt{Th}(v^-)\big)$.*
- *In poorman bidding, we have $\texttt{Th}(v) = \texttt{Th}(v^+)/(1 + \texttt{Th}(v^+) - \texttt{Th}(v^-))$.*
- *In taxman bidding with parameter τ, we have $\texttt{Th}(v) = \big(\texttt{Th}(v^-) + \texttt{Th}(v^+) - \tau \cdot \texttt{Th}(v^-)\big)/\big(2 - \tau \cdot (1 + \texttt{Th}(v^-) - \texttt{Th}(v^+))\big)$.*

Moreover, only double-reachability Richman-bidding games exhibit the following probabilistic connection: for every vertex v, we have $\texttt{Th}(v) = 1 - val\big(\texttt{RT}(\mathcal{G}), v\big)$. Thus, for games played on finite graphs, the threshold ratios are all rational numbers. However, threshold ratios with poorman-bidding need not be rational in finite games.

The equivalence between double-reachability bidding games and reachability games with Richman- and poorman-bidding is shown in [8] and [9]. The following proposition is the key component for a similar equivalence using taxman bidding.

Lemma 6. *Consider a taxman reachability game $\mathcal{G}$. Suppose that every vertex in $\mathcal{G}$ has a path to the target of Player 1. Then, for any taxman parameter, Player 1 wins from every vertex with any positive initial budget. Thus, for every vertex v, we have $\texttt{Th}(v) = 0$.*

Proof. Let $\mathcal{G} = \langle V, E, t\rangle$, where $n = |V| - 1$. Suppose the game starts from a vertex v. Let $\epsilon > 0$ be the initial budget of Player 1 and $1 - \epsilon$ be Player 2's budget. Since there is a path from v to Player 1's target, there is a path of length at most n. Thus, if Player 1 wins n consecutive biddings, he wins the game. Intuitively, Player 1 carefully chooses n increasing bids such that if Player 2 wins one of these bids, Player 1's ratio increases by a constant over his initial budget. By repeatedly playing according to such a strategy, Player 1 guarantees that his ratio increases and will eventually allow him to win n biddings in a row. Formally, if $\tau = 0$, then $\mathcal{G}$ is a Richman game and the proof of the lemma can be found in [8, Proposition 7]. Otherwise, pick a sufficiently large $r \in \mathbb{N}$ such that $\tau > \frac{2}{r-1}$ and $r \geq 3$. Fix $0 < m < \frac{\epsilon}{r^n}$. Player 1 proceeds as follows: after winning i times, for $0 \leq i$, he bids $m \cdot r^i$ and, upon winning the bidding, he moves towards t along any shortest path. Since $m + mr + \cdots + mr^{n-1} < mr^n < \epsilon$, Player 1 has sufficient budget to win n consecutive biddings. If Player 2 does not win any of the first n biddings, Player 1 wins the game.

On the other hand, if Player 2 wins the k-th bidding with $1 \leq k \leq n$, we show that Player 1's ratio increases by a fixed amount $b = \frac{mr}{(1-\epsilon)(r-1)} > 0$. Recall that the players' initial budgets are ϵ for Player 1 and $1 - \epsilon$ for Player 2. If Player 2 wins the k-th bidding with $1 \leq k \leq n$, then Player 1's new ratio is

$$\begin{aligned}
\frac{\epsilon - m - mr - \cdots - mr^{k-1} + \tau mr^k}{1 - \epsilon + \tau m + \tau mr + \cdots + \tau mr^{k-1} - mr^k} &= \frac{\epsilon - m\frac{r^k-1}{r-1} + \tau mr^k}{1 - \epsilon + \tau m\frac{r^k-1}{r-1} - mr^k} \\
&> \frac{\epsilon + mr^k(\tau - \frac{1}{r-1})}{1 - \epsilon - mr^k(1 - \frac{\tau}{r-1})} \\
&> \frac{\epsilon + mr^k(\tau - \frac{1}{r-1})}{1 - \epsilon} \\
&\overset{\tau > \frac{2}{r-1}}{>} \frac{\epsilon}{1-\epsilon} + \frac{mr^k}{(1-\epsilon)(r-1)} \\
&\geq \frac{\epsilon}{1-\epsilon} + \frac{mr}{(1-\epsilon)(r-1)},
\end{aligned}$$

Thus, the ratio increases by a fixed amount $b = \frac{mr}{(1-\epsilon)(r-1)} > 0$. Let ϵ_1 be the new (normalized) ratio of Player 1. Since $0 < m < \frac{\epsilon}{r^n} < \frac{\epsilon_1}{r^n}$, Player 1 can repeat the same process and again either win the game in at most n steps or increase his budget ratio by at least b. Note that $\frac{mr}{(1-\epsilon)(r-1)}$ is an increasing function of ϵ. Proceeding like this, eventually either Player 1 wins the game, or his normalized budget exceeds $1 - 2^{-n}$, in which case he can win n consecutive biddings by bidding $2^{-n}, 2^{-n+1}, \ldots, 2^{-1}$. □

The following corollary shows the equivalence between reachability and double-reachability taxman games.

Corollary 7. *Consider a reachability taxman game $\mathcal{G} = \langle V, E, t\rangle$. Let $S \subseteq V$ be the set of vertices that have no path to t. Let $T \subseteq V$ be a set of vertices such that $t \in T$ and every $u \in T$ has a path to t and no path to a vertex in S. Then, for every $v \in T$, we have $\mathtt{Th}(v) = 0$, for every $v \in S$, we have $\mathtt{Th}(v) = 1$. Let $\mathcal{G}'$ be a double-reachability taxman game that is obtained from $\mathcal{G}$ by merging the vertices in S and T into two targets t_1 and t_2 for Players 1 and 2, respectively. Then, for every $v \in (V \setminus (S \cup T))$, the threshold of v in $\mathcal{G}$ equals the threshold of v in $\mathcal{G}'$.*

A simple reduction from parity bidding games to reachability games was shown for Richman and poorman bidding games. A reduction for taxman bidding follows the same idea, and we describe it below for completeness. Intuitively, we show that in each BSCC, one of the players wins with any initial ratio. Then, we play a reachability game on the rest of the graph in which each player tries to force the game into a BSCC that is winning for him.

Theorem 8. *Parity taxman games are linearly reducible to taxman reachability games. Specifically, threshold ratios exist in parity taxman games.*

Proof. Consider a BSCC S. Let t be a vertex with highest parity index. We claim if the index of t is odd, then Player 1 wins from every vertex in S with any positive initial ratio. Dually, Player 2 wins when the index of t is odd. We prove for the odd case and the even case is dual.

We show how Lemma 6 implies that Player 1 wins from every vertex in S with any positive initial ratio. The construction of a winning strategy proceeds in two steps. First, we claim that with any positive ratio, Player 1 can force the game to t once. Indeed, consider the reachability game that is played on S in which Player 1's target is t and Player 2 has no target. Lemma 6 shows a winning strategy for Player 1 for any initial ratio and any taxman parameter. Second, to force the game to t infinitely often with an initial ratio of $\epsilon > 0$, Player 1 splits his budget into parts $\epsilon_1, \epsilon_2, \ldots$ and uses ϵ_i, for $i \geq 1$, to force the game to visit t for the i-th time. Since the highest parity index is obtained in t and it is odd, the resulting play is winning for Player 1.

Let S and T denote the set of vertices in BSCCs that are winning for Player 1 and Player 2, respectively. The argument above shows that the thresholds of the vertices in $S \cup T$ are either 0 or 1 depending on the highest parity index in the BSCC. We show that threshold ratios exist in all other vertices by reducing the game to a double-reachability game. We replace S with a target t_1 for Player 1 and T with a target t_2 for Player 2.

Consider a vertex v in the resulting game with a threshold ratio $\mathtt{Th}(v)$. Suppose Player i's budget is B_i, for $i \in \{1, 2\}$, and that Player 1's ratio is $r > \mathtt{Th}(v)$, and the case where Player 1's ratio is $r < \mathtt{Th}(v)$ is dual. Player 1 wins as follows. Let $\epsilon > 0$ such that $r' = (B_1 - \epsilon)/(B_1 - \epsilon + B_2) > \mathtt{Th}(v)$. Player 1 uses a ratio of r' to draw the game to a vertex in S. Once the game reaches S, he uses the strategy above with a budget of ϵ to win the game. □

4. Mean-payoff taxman games

This section consists of our main technical contribution. We start by showing a complete classification of the value in strongly-connected mean-payoff taxman games depending on the taxman parameter τ and the initial ratio. We then extend the solution to general games, where the solution to strongly-connected games constitutes the main ingredient in the solution of the general case.

4.1. Strongly-connected mean-payoff taxman games

We start by formally defining the value of a strongly-connected mean-payoff game. Lemma 6 implies that in a strongly-connected game, a player can draw the game from every vertex to any other vertex with any positive initial budget. Since mean-payoff objectives are prefix independent, it follows that the vertex from which the game starts does not matter. Indeed, if the game starts at a vertex v with Max having initial ratio $r + \epsilon$, then Max can use $\epsilon/2$ of his budget to draw the game to a vertex u and continue as if he starts the game with initial ratio $r + \epsilon/2$.

Definition 9 ***(Mean-payoff value)***. Consider a strongly-connected mean-payoff game $\mathcal{G}$ and a ratio $r \in (0, 1)$ and a taxman parameter $\tau \in [0, 1]$. The mean-payoff value of $\mathcal{G}$ w.r.t. r and τ, is a value $c \in \mathbb{R}$ such that for every $\epsilon > 0$

- if Min's initial ratio is greater than $(1 - r)$, then he has a strategy that guarantees that the payoff is at most $c + \epsilon$, and
- if Max's initial ratio is greater than r, then he has a strategy that guarantees that the payoff is greater than $c - \epsilon$.

The following theorem, which we prove in the next two sections, summarizes the properties of mean-payoff taxman games.

Theorem 10. *Consider a strongly-connected mean-payoff taxman game $\mathcal{G}$ with taxman parameter $\tau \in [0, 1]$ and an initial ratio $r \in (0, 1)$. The value of $\mathcal{G}$ w.r.t. τ and r equals the value of the random-turn game $\mathtt{RT}(\mathcal{G}, \frac{r+\tau(1-r)}{1+\tau})$.*

Remark 11. An observation made in [11] shows a simplified proof of Lemma 6 based on Theorem 10. We keep the direct proof since we believe the explicit construction has merit. For ease of presentation, we describe the construction for Büchi games and the extension to parity games is immediate. We construct a mean-payoff game $\mathcal{G}'$ by setting the weight of every accepting vertex to be 1 and every rejecting vertex to be 0. It is shown in [11] that for every $p \in (0, 1)$, we have $\texttt{MP}\big(\texttt{RT}(\mathcal{G}', p)\big) > 0$. Suppose Player 1 plays according to a strategy for Max in $\mathcal{G}'$ that guarantees a positive payoff. We claim that Player 1 wins. Indeed, we associate Player 2 with Min and no matter which strategy he chooses, the infinite play visits an accepting vertex infinitely often otherwise the payoff would not be positive in $\mathcal{G}'$.

We show that in order to prove Theorem 10, it suffices to prove the following intermediate lemma.

Lemma 12. *Consider a strongly-connected mean-payoff taxman game $\mathcal{G}$, a taxman parameter τ, and an initial ratio $r \in (0, 1)$ such that $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, \frac{r+\tau(1-r)}{1+\tau})\big) = 0$. Then, for every $\epsilon > 0$ Max has a strategy that guarantees that no matter how Min plays, the payoff is greater than $-\epsilon$.*

Proof: Lemma 12 implies Theorem 10. We may assume that we have $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, \frac{r+\tau(1-r)}{1+\tau})\big) = 0$ since we can decrease all weights by $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, \frac{r+\tau(1-r)}{1+\tau})\big)$.

Intuitively, Lemma 12 constructs a Max strategy in $\mathcal{G}$ that guarantees a payoff of at least $-\epsilon$, for $\epsilon > 0$. Recall that the definition of payoff uses lim inf and thus favors Min. In other words, the construction in Lemma 12 handles the more challenging case. We use it to obtain a strategy for Min that guarantees a payoff of at most ϵ. Technically, we consider the "dual" game $\mathcal{G}'$ that is obtained from $\mathcal{G}$ by multiplying all weights by -1. We associate Min with a ratio of r in $\mathcal{G}$ with Max with a ratio of $1 - r$, and vice versa. We show that the value of $\mathcal{G}'$ w.r.t. $1 - r$ and τ is 0, thus a Max strategy that guarantees a payoff of at least $-\epsilon$ in $\mathcal{G}'$ translates to a Min strategy guaranteeing a payoff of at most ϵ in $\mathcal{G}$.

Formally, we claim that $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, x)\big) = -\texttt{MP}\big(\texttt{RT}(\mathcal{G}', 1 - x)\big)$ for any $x \in (0, 1)$. Let $m = \texttt{MP}\big(\texttt{RT}(\mathcal{G}, x)\big)$. We prove that $\texttt{MP}\big(\texttt{RT}(\mathcal{G}', 1 - x)\big) \le -m$ and the other inequality follows by symmetry. Stochastic mean-payoff games admit positional optimal strategies. Every two positional strategies f and g induce a Markov chain on the underlying game graph, and the resulting mean-payoff is simply the weighted sum $\sum_{v \in V(G)} \pi_{f,g}(v) \cdot w(v)$, where $\pi_{f,g}$ is the stationary distribution of this Markov chain. Thus, any pair of positional strategies induce a mean-payoff that is independent of whether we use lim inf or lim sup in the definition. Note that in each turn, Min is chosen in $\texttt{RT}(\mathcal{G}', 1 - x)$ with probability x. To guarantee a payoff of at most $-m$ in $\texttt{RT}(\mathcal{G}', 1 - x)$, Min follows an optimal positional strategy for Max in $\mathcal{G}$.

Suppose Max's ratio in $\mathcal{G}$ is r, thus Min's ratio in $\mathcal{G}$ is $1 - r$. According to Lemma 12, the random-turn game that corresponds to $\mathcal{G}$ uses a coin with bias $p_{\tau,r} = \frac{r+\tau(1-r)}{1+\tau}$ towards Max and $1 - p_{\tau,r}$ towards Min. In the corresponding game $\mathcal{G}'$ the budget of Max is $1 - r$ and the budget of Min is r. Thus, according to Lemma 12 the coin chooses Max in $\mathcal{G}'$ with bias $p_{\tau,(1-r)}$. A simple calculation reveals that $1 - p_{\tau,r} = p_{\tau,(1-r)}$, thus $\texttt{MP}\big(\texttt{RT}(\mathcal{G}', p_{\tau,1-r})\big) = -\texttt{MP}\big(\texttt{RT}(\mathcal{G}, 1 - p_{\tau,1-r})\big) = -\texttt{MP}\big(\texttt{RT}(\mathcal{G}, p_{\tau,r})\big) = 0$. Hence, using a strategy for Max in $\mathcal{G}'$ that guarantees a payoff that is greater than $-\epsilon$ can be used by Min to guarantee a payoff in $\mathcal{G}$ that is smaller than ϵ. □

4.2. The importance of moving

The first part of the construction of an optimal strategy for Max as in Lemma 12 is to assign, to each vertex $v \in V$, a *strength*, denoted $\text{St}(v)$, where $\text{St}(v) \in \mathbb{Q}_{\ge 0}$. Intuitively, if $\text{St}(v) > \text{St}(u)$, for $u, v \in V$, it is more important for Max to move in v than it is in u. We slightly generalize the construction in [9], which in turn is a generalization of the construction in [8]. The construction relies on the concept of *potentials*, which is a well-known concept in stochastic games (see [10]) and was originally defined in the context of the strategy iteration algorithm [12]. For completeness, we present the definitions below.

Consider a strongly-connected mean-payoff game $\mathcal{G}$, and let $p \in [0, 1]$. Let f and g be two optimal positional strategies in $\texttt{RT}(\mathcal{G}, p)$, for Min and Max, respectively. For a vertex $v \in V$, let $v^-, v^+ \in V$ be such that Max proceeds from v to v^+ according to g and Min proceeds from v to v^- according to f. It follows from the fact that $\mathcal{G}$ is strongly-connected, that the mean-payoff value in all vertices in $\texttt{RT}(\mathcal{G}, p)$ is the same and we denote it by $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, p)\big)$. We denote the potential of v by $\text{Pot}^p(v)$ and the strength of v by $\text{St}^p(v)$, and we define them as follows.

$$\begin{aligned} &\text{Pot}^p(v) = p \cdot \text{Pot}^p(v^+) + (1 - p) \cdot \text{Pot}^p(v^-) + w(v) - \texttt{MP}\big(\texttt{RT}(\mathcal{G}, p)\big) \text{ and} \\ &\text{St}^p(v) = p \cdot (1 - p) \cdot \big(\text{Pot}^p(v^+) - \text{Pot}^p(v^-)\big) \end{aligned} \tag{1}$$

We assume v^- and v^+ satisfy $\text{Pot}^p(v^-) \le \text{Pot}^p(v') \le \text{Pot}^p(v^+)$, for every $v' \in N(v)$. An optimal strategy that ensures all of these constraints can be found, for example, using the strategy iteration algorithm. Note that $\text{St}(v) \ge 0$, for every $v \in V$.

We suggest an intuitive way to read the definition of potentials. Consider a weighted Markov chain in which each vertex v has two neighbors and the probability of proceeding from v to v^+ and v^- is respectively p and $1 - p$. Then, the potential of v roughly coincides with the expression to compute the expected energy in a path from v to the target.

Consider a finite path $\pi = v_1, \ldots, v_n$ in $\mathcal{G}$. We intuitively think of π as a play, where for every $1 \le i < n$, the bid of Max in v_i is $\text{St}(v_i)$ and he moves to v_i^+ upon winning. Thus, if $v_{i+1} = v_i^+$, we say that Max won in v_i, and if $v_{i+1} \ne v_i^+$, we say that Max lost in v_i. Let $W(\pi)$ and $L(\pi)$ respectively be the indices in which Max wins and loses in π. We call Max wins *investments* and Max loses *gains*, where intuitively he *invests* in increasing the energy and *gains* a higher ratio of the budget whenever the energy decreases. Let $G(\pi)$ and $I(\pi)$ be the sum of gains and investments in π, respectively, thus $G(\pi) = \sum_{i\in L(\pi)} \mathrm{St}(v_i)$ and $I(\pi) = \sum_{i\in W(\pi)} \mathrm{St}(v_i)$. Recall that the energy of π is $E(\pi) = \sum_{1\le i<n} w(v_i)$. The following lemma, which generalizes a similar lemma in [9], connects the strength with the change in energy.

Lemma 13. *Consider a strongly-connected game $\mathcal{G}$, and $p \in [0,1]$. For every finite path $\pi = u_1,\dots,u_n$ in $\mathcal{G}$, we have $Pot^p(u_1) - Pot^p(u_n) + (n-1)\cdot MP\big(RT(\mathcal{G},p)\big) \le E(\pi) + G(\pi)/(1-p) - I(\pi)/p$. In particular, when $p = \nu/(\mu+\nu)$ for $\nu,\mu>0$, there is a constant $P = \min_u Pot^p(u) - \max_u Pot^p(u)$ such that $\frac{\nu\cdot\mu}{\nu+\mu}\cdot\big(E(\pi) - P - (n-1)\cdot MP\big(RT(\mathcal{G},\frac{\nu}{\mu+\nu})\big)\big) \ge \mu\cdot I(\pi) - \nu\cdot G(\pi)$.*

Proof. The proof is by induction on the length of the path. Recall that the weight of the last vertex of a finite path does not contribute to the energy, thus for a path $\pi = u_1,\dots,u_n,u_{n+1}$, we have $E(\pi) = \sum_{1\le i\le n} w(u_i)$.

For the base case, when $\pi = u_1$, we have $n=1$. Since no biddings take place and the weight of the last (first) vertex does not contribute to the energy, we have $I(\pi) = G(\pi) = E(\pi) = 0$. Thus, both sides of the equation are 0. Suppose the claim is true for all paths of length n and we prove it for a path $\pi = u_1,\dots,u_{n+1}$. We consider the case when Max wins in u_1 thus $u_2 = u_1^+$. The case when Min wins in u_1 is proved similarly. Let π' be the part of path π starting in u_2. Since Max wins the first bidding, we have $G(\pi') = G(\pi)$, $I(\pi') = I(\pi) + \mathrm{St}^p(u_1)$. Hence, by the induction hypothesis we have

$$
\begin{aligned}
E(\pi) + G(\pi)/(1-p) - I(\pi)/p &\ge E(\pi') + G(\pi')/(1-p) - I(\pi')/p + w(u_1) - \mathrm{St}^p(u_1)/p \\
&\ge \mathrm{Pot}^p(u_1^+) - \mathrm{Pot}^p(u_{n+1}) + (n-1)\cdot \mathtt{MP}\big(\mathtt{RT}(\mathcal{G},p)\big) + w(u_1) - \mathrm{St}^p(u_1)/p \\
&= \mathrm{Pot}^p(u_1^+) - \mathrm{Pot}^p(u_{n+1}) + (n-1)\cdot \mathtt{MP}\big(\mathtt{RT}(\mathcal{G},p)\big) + w(u_1) - (1-p)\cdot(\mathrm{Pot}^p(u_1^+) - \mathrm{Pot}^p(u_1^-)) \\
&= p\cdot \mathrm{Pot}^p(u_1^+) + (1-p)\cdot \mathrm{Pot}^p(u_1^-) + w(u_1) - \mathrm{Pot}^p(u_{n+1}) + (n-1)\cdot \mathtt{MP}\big(\mathtt{RT}(\mathcal{G},p)\big) \\
&= \mathrm{Pot}^p(u_1) - \mathrm{Pot}^p(u_{n+1}) + n\cdot \mathtt{MP}\big(\mathtt{RT}(\mathcal{G},p)\big). \quad \square
\end{aligned}
$$

4.3. Normalizing the bids

Whenever the game reaches a vertex v, Max's bid is obtained by carefully normalizing the strength of v. More formally, assuming an initial ratio r, in v, Max bids $r\cdot(1-r)\cdot \mathrm{St}(v)\cdot\beta_x$, where β_x is the normalization factor and $x\in\mathbb{R}_{\ge 1}$. In this section we show how to choose the normalization factor. We associate with every $x\ge 1$, two numbers: a ratio r_x and β_x both in $(0,1)$. We think of $\{r_x\}_{x\ge 1}$ as a sequence and a play gives rise to a walk on the sequence, which corresponds to the changes in energy in the bidding game. Suppose the walk is at $x\ge 1$, then Max bids $r\cdot(1-r)\cdot \mathrm{St}(v)\cdot\beta_x$. If Max wins, we take a step up on the sequence, modeling the increase of energy, and when Min wins, we talk a step down. The size of the step depends on the strength of v. We formalize the properties of the sequence in the following lemma and formally define Max's strategy after it.

Lemma 14. *Consider a game $\mathcal{G}$, a finite set of non-negative strengths $S\subseteq\mathbb{R}_{\ge 0}$, a ratio $r\in(0,1)$, and a taxman parameter $\tau\in[0,1]$. For every $K > \frac{\tau r^2 + r(1-r)}{\tau(1-r)^2 + r(1-r)}$ there exist sequences $(r_x)_{x\ge 1}$ and $(\beta_x)_{x\ge 1}$ with the following properties.*

1. *Max's bid does not exceed his budget, thus, for each position $x\in\mathbb{R}_{\ge 1}$ and strength $s\in S$, we have $\beta_x\cdot s\cdot r\cdot(r-1) < r_x$.*
2. *Min cannot force the game beyond position 1, thus for every $s\in S\setminus\{0\}$ and $1\le x<1+rs$, we have $\beta_x\cdot s\cdot r\cdot(r-1) > 1-r_x$.*
3. *The ratios tend to r from above, thus for every $x\in\mathbb{R}_{\ge 1}$, we have $r_x\ge r$, and $\lim_{x\to\infty} r_x = r$.*
4. *No matter who wins a bidding, Max's ratio can only improve. Thus, in case of winning and in case of losing, we respectively have*

$$\frac{r_x - \beta_x\cdot s\cdot r\cdot(r-1)}{1-(1-\tau)\cdot\beta_x\cdot s\cdot r\cdot(r-1)} \ge r_{x+(1-r)\cdot K\cdot s}$$

and

$$\frac{r_x + \tau\cdot\beta_x\cdot s\cdot r\cdot(r-1)}{1-(1-\tau)\cdot\beta_x\cdot s\cdot r\cdot(r-1)} \ge r_{x-s\cdot r}.$$

We first show how Lemma 14 implies Theorem 10.

Proof: Lemma 14 implies Lemma 12. Fix $\epsilon>0$, we construct a strategy for Max guaranteeing a payoff greater than $-\epsilon$, as required. Observe that

$$\frac{r}{r+(1-r)\frac{\tau r^2+r(1-r)}{\tau(1-r)^2+r(1-r)}} = \frac{r(\tau(1-r)+r)}{\tau r(1-r)+r^2+\tau r^2+r(1-r)} = \frac{r+\tau(1-r)}{1+\tau}.$$

Thus, since by assumption $\mathtt{MP}\big(\mathtt{RT}(\mathcal{G}, \frac{r+\tau(1-r)}{1+\tau})\big) = 0$ and $\mathtt{MP}\big(\mathtt{RT}(\mathcal{G}, p)\big)$ is a continuous function in $p \in [0, 1]$ [13,14], we can pick $K > \frac{\tau r^2 + r(1-r)}{\tau(1-r)^2 + r(1-r)}$ such that $\mathtt{MP}\big(\mathtt{RT}(\mathcal{G}, \frac{r}{r+(1-r)K})\big) > -\epsilon$.

We now describe Max's strategy. We think of the change in Max's ratio as a walk on $\mathbb{R}_{\geq 1}$. Each position $x \in \mathbb{R}_{\geq 1}$ is associated with a ratio r_x. The walk starts in a position x_0 such that Max's initial ratio is at least r_{x_0}. Note that it is possible to pick such x_0 by Point 3 of Lemma 14 since $r_1 = 1$, $\lim_{x\to\infty} r_x = r$, r_x is continuous in x and the initial budget of Max exceeds r. Let $\nu = r$ and $\mu = K(1-r)$. Suppose the token is placed on a vertex $v \in V$. Then, Max's bid is $r \cdot (1-r) \cdot \beta_x \cdot \mathrm{St}(v)$ (such bid is legal by Point 1 of Lemma 14), where the ratios of Max and Min are normalized to sum up to 1, and he proceeds to v^+ upon winning. If Max wins, the walk proceeds up $\mu \cdot \mathrm{St}(v)$ steps to $x + \mu\mathrm{St}(v)$, and if he loses, the walk proceeds down to $x - \nu\mathrm{St}(v)$. Point 4 of Lemma 14 ensures that the budget ratio of Max is at least r_x. We show that this strategy of Max ensures payoff of at least $-\epsilon$, which proves Lemma 12.

Suppose Min fixes some strategy and let $\pi = v_1, \ldots, v_n$ be a finite prefix of the play that is generated by the two strategies. Suppose the walk following π reaches $x \in \mathbb{R}$. Point 2 of Lemma 14 shows that with Max using this strategy, Min cannot win a bid if the win would force the walk below 1, therefore the walk always stays above 1 and $x \geq 1$. Then, using the terminology of the previous section, we have $x = x_0 - G(\pi) \cdot \nu + I(\pi) \cdot \mu \geq 1$. Combining with Lemma 13, we get $\frac{\nu+\mu}{\nu\cdot\mu}(1 - x_0) + P + (n-1) \cdot \mathtt{MP}\big(\mathtt{RT}(\mathcal{G}, \frac{\nu}{\nu+\mu})\big) \leq E(\pi)$. Thus, dividing both sides by n and letting $n \to \infty$, since x_0 and P are constants depending only on K we conclude that this strategy guarantees a payoff of at least $\mathtt{MP}\big(\mathtt{RT}(\mathcal{G}, \frac{\nu}{\nu+\mu})\big) > -\epsilon$. □

We continue to prove Lemma 14.

Proof of Lemma 14. Note that $\frac{\tau r^2 + r(1-r)}{\tau(1-r)^2 + r(1-r)}$ is well-defined for $r \in (0, 1)$. Fix $\tau \in [0, 1]$ and $r \in (0, 1)$. Let $K > \frac{\tau r^2 + r(1-r)}{\tau(1-r)^2 + r(1-r)}$. We first show that we may replace inequalities in Point 4 with inequalities that together with Points 1-3 impose stricter conditions on $(r_x)_{x\geq 1}$ and $(\beta_x)_{x\geq 1}$, but which will make solving these conditions easier. By multiplying both sides of the inequalities in Point 4 by fraction denominators, we can see that the two inequalities in Point 4 are equivalent to:

$$\begin{aligned} r_{x-rs} - r_x &\leq \tau r(1-r)\beta_x s + (1-\tau) r(1-r)\beta_x s r_{x-rs}, \\ r_x - r_{x+K(1-r)s} &\geq r(1-r)\beta_x s - (1-\tau) r(1-r)\beta_x s r_{x+K(1-r)s}. \end{aligned}$$

If the sequences $(r_x)_{x\geq 1}$ and $(\beta_x)_{x\geq 1}$ satisfy Points 1-3 in Lemma 14, then by Point 3 $r_x \geq r$ for each $x \in \mathbb{R}_{\geq 1}$. Therefore we may substitute $r_{x-rs} \geq r$ and $r_{x+K(1-r)s} \geq r$ in the above inequalities to obtain a stricter condition. Thus, it suffices to show that there exist sequences $(r_x)_{x\geq 1}$ and $(\beta_x)_{x\geq 1}$ that satisfy Points 1-3 in Lemma 14 and the following inequalities:

$$\begin{aligned} r_{x-rs} - r_x &\leq \tau r(1-r)\beta_x s + (1-\tau) r(1-r)\beta_x s r, \\ r_x - r_{x+K(1-r)s} &\geq r(1-r)\beta_x s - (1-\tau) r(1-r)\beta_x s r, \end{aligned}$$

which are equivalent to

$$\begin{aligned} r_{x-rs} - r_x &\leq r(1-r)\beta_x s[\tau + (1-\tau)r], \\ r_x - r_{x+K(1-r)s} &\geq r(1-r)\beta_x s[1 - (1-\tau)r]. \end{aligned} \tag{2}$$

We seek $(r_x)_{x\geq 1}$ and $(\beta_x)_{x\geq 1}$ in the form $r_x = \gamma^{x-1} + (1-\gamma^{x-1})r$ and $\beta_x = \beta\gamma^{x-1}$ for some $\gamma, \beta \in (0, 1)$. Hence, in the rest of this proof we show that it is possible to find $\gamma, \beta \in (0, 1)$ that ensure all 4 conditions are satisfied.

For Point 1 to hold, a simple computation shows that it suffices to have $\beta < 1/(\max_{s\in S} s)$. Point 3 trivially holds for any $\gamma, \beta \in (0, 1)$ since $r_1 = 1$, $\lim_{x\to\infty} r_x = r$ and r_x is decreasing in x.

We now find conditions for $\gamma, \beta \in (0, 1)$ to satisfy the inequalities in (2) for any $s \in S$. Substituting r_x and β_x in terms of γ and β, the inequalities in (2) are reduced to

$$\begin{aligned} r_{x-rs} - r_x &= \gamma^{x-1}(\gamma^{-rs} - 1)(1-r) \overset{?}{\leq} \beta\gamma^{x-1} r(1-r)s[\tau + (1-\tau)r], \\ r_x - r_{x+K(1-r)s} &= \gamma^{x-1}(1 - \gamma^{K(1-r)s})(1-r) \overset{?}{\geq} \beta\gamma^{x-1} r(1-r)s[1 - (1-\tau)r]. \end{aligned}$$

First, when $s = 0$, both sides of both inequalities are equal to 0 so both inequalities clearly hold. Recall that S is a finite set of non-negative strengths. Thus, when $s > 0$, it takes values in $0 < s_1 \leq \ldots \leq s_n$, and a simple computation shows that the above inequalities are equivalent to

$$\begin{aligned} \gamma &\geq \big(1 + \beta rs[\tau + (1-\tau)r]\big)^{-\frac{1}{rs}}, \\ \gamma &\leq \big(1 - \beta rs[1 - (1-\tau)r]\big)^{\frac{1}{K(1-r)s}}. \end{aligned} \tag{3}$$

Since both of these expressions are in $(0, 1)$, to conclude that $\gamma, \beta \in (0, 1)$ exist, it suffices to show that there is some $\beta \in (0, 1)$ such that

$$\max_{s\in\{s_1,\dots,s_n\}} \big(1+\beta rs[\tau+(1-\tau)r]\big)^{-\frac{1}{rs}} \le \min_{s\in\{s_1,\dots,s_n\}} \big(1-\beta rs[1-(1-\tau)r]\big)^{\frac{1}{K(1-r)s}}. \tag{4}$$

Note that the LHS of (4) is monotonically increasing in $s>0$ whereas the RHS is monotonically decreasing in $s>0$, therefore it suffices to find $\beta\in(0,1)$ for which

$$\big(1+\beta rs_n[\tau+(1-\tau)r]\big)^{-\frac{1}{rs_n}} \le \big(1-\beta rs_1[1-(1-\tau)r]\big)^{\frac{1}{K(1-r)s_1}}. \tag{5}$$

By Taylor's theorem $(1+y)^\alpha = 1+\alpha y+O(y^2)$, so Taylor expanding both sides of (5) in $\beta>0$ we get

$$\big(1+\beta rs_n[\tau+(1-\tau)r]\big)^{-\frac{1}{rs_n}} = 1-\beta[\tau+(1-\tau)r]+O(\beta^2),$$
$$\big(1-\beta rs_1[1-(1-\tau)r]\big)^{\frac{1}{K(1-r)s_1}} = 1-\beta\frac{r}{K(1-r)}[1-(1-\tau)r]+O(\beta^2).$$

Therefore, if we show that $[\tau+(1-\tau)r] > \frac{r}{K(1-r)}[1-(1-\tau)r]$, the linear coefficient of β on the LHS of (5) will be strictly smaller than the linear coefficient of β on the RHS. Thus, for sufficiently small $\beta>0$, (5) will hold, which concludes the proof of the lemma. This condition is equivalent to

$$K > \frac{r[1-(1-\tau)r]}{(1-r)[\tau+(1-\tau)r]} = \frac{r[\tau r+(1-r)]}{(1-r)[\tau(1-r)+r]} = \frac{\tau r^2+r(1-r)}{\tau(1-r)^2+r(1-r)},$$

which is true by assumption. Thus, Points 1, 3, and 4 hold for sufficiently small value of β (with $\beta<1/(\max_{s\in S} s)$ to ensure Point 1) and the value of γ between the values defined in the inequality (5).

We conclude the proof by showing that Point 2 holds for any $\gamma,\beta\in(0,1)$. Let $s\in S\setminus\{0\}$ and $1\le x<1+rs$. Intuitively, if Min wins the bidding, we reach a position that is less than 1. We show that $1-r_x < sr(1-r)\beta_x$, therefore proving that Min has insufficient budget to win this bid. Taking $(\gamma_x)_{x\ge 1}$ and $(\beta_x)_{x\ge 1}$ as in the above, we have $1-r_x=(1-\gamma^{x-1})(1-r)$ and $\beta_x=\beta\gamma^{x-1}$. Hence it suffices to prove that $\gamma^{x-1} > \frac{1}{1+sr\beta}$. As $x-1<sr$ and $\gamma\in(0,1)$, we have $\gamma^{x-1}>\gamma^{sr}$. On the other hand, we established the first inequality in (3), thus as $[\tau+(1-\tau)r]\le 1$ and $sr\beta\ge 0$, we conclude $\gamma^{sr}\ge \frac{1}{1+sr\beta[\tau+(1-\tau)r]} \ge \frac{1}{1+sr\beta}$. □

4.4. General mean-payoff taxman games

In this section we extend the solution to strongly-connected games to games played on general graphs. Intuitively, the question we consider regards the necessary and sufficient budget ratio to guarantee a given target payoff from a given vertex. Formally, we consider WLog a target payoff of 0 since if the target is $c\neq 0$, we can always decrease c from all weights in the game. Recall that for a vertex v in a (general) mean-payoff game, we define $\mathtt{Th}(v)=t$ if for every $\epsilon>0$, Max can guarantee a payoff of at least $-\epsilon$ with a ratio greater than t and Min can guarantee a payoff of at most ϵ with an initial ratio greater than $1-t$.

The proof technique is similar to the reduction in parity games. We first consider the BSCCs of the graph, whose solution follows from Theorem 10 as follows. In a BSCC S, the required ratio to ensure a payoff of 0 is $r(S)=\inf_r \mathtt{MP}\big(\mathtt{RT}(S,\frac{r+\tau(1-r)}{1+\tau})\big)\ge 0$. When $\mathtt{MP}\big(\mathtt{RT}(S,1)\big)<0$, we set $r(S)=1$, which vacuously states that only when Max's ratio is greater than 1, he can ensure a payoff of 0. Note that $r(S)$ might be 0 in which case even if Min wins all biddings in S, the payoff is non-negative. This occurs when all cycles in S have non-negative weight. In between, i.e., $r(S)\in(0,1)$, Theorem 10 implies that when Max's ratio exceeds $r(S)$ he can guarantee a payoff of at least $-\epsilon$, for every $\epsilon>0$. Dually, when Max's ratio is less than r, Min can guarantee a payoff of at most $\epsilon>0$.

To determine the threshold ratios in vertices that do not belong to BSCCs we construct a *generalized-reachability* bidding game, which is defined as follows. The game is played on an arena $\langle V,E,\{\langle t_i,r_i\rangle\}_{1\le i\le k}\rangle$, where $t_i\in V$ is a target vertex and $r_i\in[0,1]$, for $1\le i\le k$. The game ends once a target t_i is reached and Player 1 wins iff his ratio exceeds r_i. The solution to generalized-reachability taxman games requires the following slight adjustment to the solution to reachability taxman games from [4, Theorem 6]: while in reachability taxman games the ratios in the targets are 0 or 1, here targets can take numbers in the whole interval [0, 1]. The winning strategy constructed in [4] has two properties: first, it maintains the invariant that Player 1's ratio exceeds $\mathtt{Th}(v)$ whenever the game visits v. Second, similar to the proof of Lemma 6, Player 1 plays so that either a target is reached or the difference between his ratio and $\mathtt{Th}(v)$ increases, thereby guaranteeing that eventually a target is reached.

To conclude, given a mean-payoff taxman game $\mathcal{G}$ with taxman parameter τ, for every BSCC S, we find the threshold $r(S)$ as in the above. We construct a generalized-reachability game $\mathcal{G}'$ by replacing every S with a target t with the ratio $r(S)$. For a vertex v, let $\mathtt{Th}(v)$ be the threshold ratio in $\mathcal{G}'$. Suppose $\mathcal{G}$ starts at v with Max's ratio exceeding $\mathtt{Th}(v)$ and the claim for Min is dual. Max guarantees a payoff greater than $-\epsilon$ by following his winning strategy in $\mathcal{G}'$, which ensures that the game reaches some BSCC S with a ratio greater than $r(S)$. Once S is reached, Max turns to play according to a strategy as in Theorem 10 to guarantee a payoff greater than $-\epsilon$. We thus have the following.

Theorem 15. *Threshold ratios exist in mean-payoff taxman games.*

5. Computational complexity

We show, for the first time, computational complexity results for taxman games. Intuitively, in qualitative games, we study the problem of finding the threshold ratios in a vertex. In mean-payoff games, we find the necessary and sufficient initial ratio with which a player can guarantee a target payoff. Stated as a decision problem:

Definition 16. The input to *THRESH* is a taxman game $\mathcal{G}$ with taxman parameter τ and a vertex v in $\mathcal{G}$, and the goal is to decide whether $\texttt{Th}(v) \geq 0.5$ in $\mathcal{G}$.

Theorem 17. *For taxman reachability, parity, and mean-payoff games THRESH is in PSPACE. For strongly-connected mean-payoff games, THRESH is in NP ∩ coNP.*

Proof. We start with strongly-connected mean-payoff games. Consider such a game $\mathcal{G}$. Then, by Theorem 10, the threshold of a vertex is at least 0.5 if and only if $\texttt{MP}\big(\texttt{RT}(\mathcal{G}, \frac{0.5+\tau\cdot 0.5}{1+\tau})\big) \geq 0$. Since random-turn games are a special case of stochastic games, the upper bound of NP∩coNP [15] of the latter follows to bidding games.

In the other cases, we describe a non-deterministic polynomial space algorithm that relies on the existential theory of the reals (ETR, for short), which is known to be in PSPACE [16]. Membership in PSPACE follows from the fact that NPSPACE=PSPACE. We describe the reduction for mean-payoff games and the solution for reachability and parity games follows similar lines. The program constitutes of two separate parts: the first corresponds to BSCCs and the second to the rest of the graph.

We start with the BSCCs. Consider a BSCC S. We define a variable $r(S)$ and verify that the threshold ratio in S is at most $r(S)$. We can check in polynomial time whether $\texttt{MP}\big(\texttt{RT}(S,1)\big) < 0$ or $\texttt{MP}(\big(\texttt{RT}(S,0)\big) > 0$. If the first is true, we add a constraint $r(S) = 1$, and if the second is true, we add a constraint $r(S) = 0$. Otherwise, we define a variable $\texttt{MP}(S)$ and verify that the threshold ratio in S is at most $r(S)$ using an input to the ETR that is based on Equation (1). We add a constraint $\texttt{MP}(S) \geq 0$. For every vertex v in S, we define a variable x_v whose assignment will intuitively be its potential. In addition, we guess two neighboring vertices v^+ and v^- of v. We add a constraint $x_{v^-} \leq x_u \leq x_{v^+}$, for every neighbor u of v. We add a constraint $x_v \geq \frac{r(S)+\tau(1-r(S))}{1+\tau} \cdot x_{v^+} + (1 - \frac{r(S)+\tau(1-r(S))}{1+\tau}) \cdot x_{v^-} + w(v) - \texttt{MP}(S)$. An assignment to the variables that satisfies the constraints must assign to each x_v the potential of v and to $\texttt{MP}(S)$, the mean-payoff value of $\texttt{RT}(S, \frac{r(S)+\tau(1-r(S))}{1+\tau})$. Since we require $\texttt{MP}(S) \geq 0$, the threshold ratio in S is at most $r(S)$, as required.

We proceed to vertices that do not belong to BSCCs. Similar to the above, we simulate the procedure described in Section 4.4 to solve generalized-reachability games using ETR. For each vertex v, we set a variable y_v, which will take the threshold value at v. For each vertex v in a BSCC S, we add a constraint $y_v = r(S)$. As before, we guess neighbors v^+ and v^- and add a constraint $y_{v^-} \leq y_u \leq y_{v^+}$ for every neighbor u of v. We add a constraint $y_v = \big(y_{v^-} + y_{v^+} - \tau \cdot y_{v^-}\big)/\big(2 - \tau \cdot (1 + y_{v^-} - y_{v^+})\big)$. We claim that an assignment to x_v that satisfies the constraints represents the threshold in v. Indeed, as described in Section 4.4, Max can maintain the invariant that his ratio exceeds x_u in each vertex u while forcing that the game eventually reaches a vertex t in a BSCC S. The invariant implies that his ratio in t exceeds $r(S)$ with which he can guarantee a payoff of at least $-\epsilon$ for every $\epsilon > 0$. □

6. Discussion

We study, for the first time, infinite-duration taxman bidding games, which span the spectrum between Richman and poorman bidding. For qualitative objectives, we show that the properties of taxman coincide with these of Richman and poorman bidding. For mean-payoff games, where Richman and poorman bidding have an elegant though surprisingly different mathematical structure, we show a complete understanding of taxman games. Our study of mean-payoff taxman games sheds light on these differences and similarities between the two bidding rules. Unlike previous proof techniques, which were ad-hoc, we expect our technique to be easier to generalize beyond taxman games, where they can be used to introduce concepts like multi-players or partial information into bidding games.

Since the conference publication of this work, extensions have been made in several directions (see [17]). We focused in this work on variants of first-price bidding; namely, only the higher bidder pays his bid. Another well-studied mechanism is *second-price* bidding; namely, the higher bidder pays the bid of the second highest bidder. Our proof techniques extend to second-price bidding games, thus we show that first- and second-price bidding games coincide. A third well-known bidding mechanism is *all-pay* bidding in which both players pay their bids. All-pay poorman bidding games were studied in [18] and shown to be technically significantly more challenging than their first-price counterparts. Inspired by the result of the present work – even though an equivalence with random-turn games does not hold for reachability taxman games, it holds for mean-payoff taxman – mean-payoff all-pay bidding games were studied in [11] and intricate equivalences with random-turn games are shown for all-pay taxman bidding games.

Other bidding games have been considered. Bidding games played on MDPs were considered in [19]. Two player non-zero-sum bidding games played on DAGs were studied in [20]. Finally, in *discrete-bidding* games [21], the granularity of the bids is restricted by assuming that the budgets are given using coins. Bidding versions of recreational games have been studied, e.g., bidding chess [22]. Infinite-duration discrete-bidding games were studied in [23], where they were shown to be a determined sub-class of concurrent games.

CRediT authorship contribution statement

Guy Avni: Conceptualization, Writing. Thomas Henzinger: Conceptualization, Writing. Đorđe Žikelić: Conceptualization, Writing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was supported in part by the Austrian Science Fund (FWF) under grants S11402-N23 (RiSE/SHiNE) and Z211-N23 (Wittgenstein Award), and from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 665385.

References

[1] K. Apt, E. Grädel, Lectures in Game Theory for Computer Scientists, Cambridge University Press, 2011.
[2] A. Pnueli, R. Rosner, On the synthesis of a reactive module, in: Proc. 16th POPL, 1989, pp. 179–190.
[3] M. Rabin, Decidability of second order theories and automata on infinite trees, Trans. Am. Math. Soc. 141 (1969) 1–35.
[4] A.J. Lazarus, D.E. Loeb, J.G. Propp, W.R. Stromquist, D.H. Ullman, Combinatorial games under auction play, Games Econ. Behav. 27 (1999) 229–264.
[5] A.J. Lazarus, D.E. Loeb, J.G. Propp, D. Ullman, Richman games, in: Games of No Chance, vol. 29, 1996, pp. 439–449.
[6] Y. Peres, O. Schramm, S. Sheffield, D.B. Wilson, Tug-of-war and the infinity Laplacian, J. Am. Math. Soc. 22 (2009) 167–210.
[7] A. Condon, The complexity of stochastic games, Inf. Comput. 96 (1992) 203–224.
[8] G. Avni, T.A. Henzinger, V. Chonev, Infinite-duration bidding games, J. ACM 66 (2019) 31.
[9] G. Avni, T.A. Henzinger, R. Ibsen-Jensen, Infinite-duration poorman-bidding games, in: Proc. 14th WINE, in: LNCS, vol. 11316, Springer, 2018, pp. 21–36.
[10] M.L. Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming, John Wiley & Sons, Inc., New York, NY, USA, 2005.
[11] G. Avni, I. Jecker, Đ. Žikelić, Infinite-duration all-pay bidding games, in: Proc. 32nd SODA, 2021, pp. 617–636.
[12] A.R. Howard, Dynamic Programming and Markov Processes, MIT Press, 1960.
[13] K. Chatterjee, Robustness of structurally equivalent concurrent parity games, in: Proc. 15th FoSSaCS, 2012, pp. 270–285.
[14] E. Solan, Continuity of the value of competitive Markov decision processes, J. Theor. Probab. 16 (2003) 831–845.
[15] U. Zwick, M. Paterson, The complexity of mean payoff games on graphs, Theor. Comput. Sci. 158 (1996) 343–359.
[16] J.F. Canny, Some algebraic and geometric computations in PSPACE, in: Proc. 20th STOC, 1988, pp. 460–467.
[17] G. Avni, T.A. Henzinger, A survey of bidding games on graphs, in: Proc. 31st CONCUR, in: LIPIcs, vol. 171, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, 2.
[18] G. Avni, R. Ibsen-Jensen, J. Tkadlec, All-pay bidding games on graphs, in: Proc. 34th AAAI, AAAI Press, 2020, pp. 1798–1805.
[19] G. Avni, T.A. Henzinger, R. Ibsen-Jensen, P. Novotný, Bidding games on Markov decision processes, in: Proc. 13th RP, 2019, pp. 1–12.
[20] R. Meir, G. Kalai, M. Tennenholtz, Bidding games and efficient allocations, Games Econ. Behav. 112 (2018) 166–193, https://doi.org/10.1016/j.geb.2018.08.005.
[21] M. Develin, S. Payne, Discrete bidding games, Electron. J. Comb. 17 (2010) R85.
[22] J. Bhatt, S. Payne, Bidding chess, Math. Intell. 31 (2009) 37–39.
[23] M. Aghajohari, G. Avni, T.A. Henzinger, Determinacy in discrete-bidding infinite-duration games, in: Proc. 30th CONCUR, in: LIPIcs, vol. 140, 2019, 20.