A Survey of Bidding Games on Graphs

Guy Avni
IST Austria, Klosterneuburg, Austria

Thomas A. Henzinger
IST Austria, Klosterneuburg, Austria

Abstract

A graph game is a two-player zero-sum game in which the players move a token throughout a graph to produce an infinite path, which determines the winner or payoff of the game. In *bidding games*, both players have budgets, and in each turn, we hold an "auction" (bidding) to determine which player moves the token. In this survey, we consider several bidding mechanisms and study their effect on the properties of the game. Specifically, bidding games, and in particular bidding games of infinite duration, have an intriguing equivalence with *random-turn* games in which in each turn, the player who moves is chosen randomly. We show how minor changes in the bidding mechanism lead to unexpected differences in the equivalence with random-turn games.

2012 ACM Subject Classification Theory of computation → Solution concepts in game theory; Theory of computation → Formal languages and automata theory

Keywords and phrases Bidding games, Richman bidding, poorman bidding, mean-payoff, parity

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2020.2

Category Invited Paper

Funding This research was supported in part by the Austrian Science Fund (FWF) under grant Z211-N23 (Wittgenstein Award).

Acknowledgements We would like to thank all our collaborators Milad Aghajohari, Ventsislav Chonev, Rasmus Ibsen-Jensen, Ismäel Jecker, Petr Novotný, Josef Tkadlec, and Đorđe Žikelić; we hope the collaboration was as fun and meaningful for you as it was for us.

1 Introduction

Two-player zero-sum games on graphs games have deep connections to foundations of logic [31] as well as numerous practical applications, e.g., verification [19], reactive synthesis [29], and reasoning about multi-agent systems [2]. The game proceeds by placing a token on one of the vertices and allowing the players to move it throughout the graph to produce an infinite trace, which determines the winner or payoff of the game.

Several "modes of moving" the token have been studied. The most well-studied mode is *turn-based* games, in which the vertices are partitioned between the players and whoever controls the vertex on which the token is placed, decides its next position. Two other modes of moving that will make appearances in this survey are *stochastic games* and *concurrent games*. Stochastic games generalize turn-based games in that some vertices are controlled by nature and from which the token moves probabilistically. In concurrent games, in each vertex, the players simultaneously select actions, and the next location is determined according to the joint (deterministic or probabilistic) action vector. For formal definitions, see [3] for example.

We study the *bidding* mode of moving. Abstractly speaking, both players have budgets, and in each turn, we hold an "auction" (bidding) to determine which player moves the token. In this survey, we consider several concrete bidding mechanisms and study the properties of the bidding games that they give rise to.

31st International Conference on Concurrency Theory (CONCUR 2020).
Editors: Igor Konnov and Laura Kovács; Article No. 2; pp. 2:1–2:21
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We emphasize that bidding is a mode of moving the token and bidding games can be studied in combination with any objective. We focus on three objectives: *reachability*, *parity*, and *mean-payoff*. We start by surveying results on reachability games that were obtained in two papers in the 1990s [24, 23]. We then turn to summarize more recent results on infinite-duration games. Our most interesting results are for mean-payoff games. In a nutshell, reachability bidding games with a specific bidding mechanism called *Richman bidding* were shown to be equivalent to a class of games called *random-turn games*, which are well-studied in their own right since the seminal paper [27]. We show a generalized equivalence between mean-payoff bidding games and random-turn games. While the equivalence for finite-duration games holds only for Richman bidding, for mean-payoff games, equivalences with random-turn games hold for a wide range of bidding mechanisms.

We keep the presentation in the survey relatively informal. For formal definitions, see the cited papers.

Bidding mechanisms

In all mechanisms that we consider, both players simultaneously submit "legal" bids that do not exceed their available budgets, and the higher bidder moves the token.

We mainly focus on two orthogonal distinctions between the mechanisms: *who pays* and *who is the recipient*. To answer the latter, two mechanisms were defined in [23]: in *Richman bidding* (named after David Richman), payments are made to the other player, and in *poorman bidding* the payments are made to the "bank" thus the money is lost. A third payment scheme called *taxman* spans the spectrum between Richman and poorman.

We make the payment schemes precise below. For $i \in \{1, 2\}$, suppose Player i's budget is B_i prior to a bidding and his bid is $b_i \in [0, B_i]$, and assume for convenience that Player 1 wins the bidding, thus $b_1 > b_2$. The budgets are updated as follows:

- **First-price:** Only the higher bidder pays.
  - **Richman:** $B'_1 = B_1 - b_1$ and $B'_2 = B_2 + b_1$.
  - **Poorman:** $B'_1 = B_1 - b_1$ and $B'_2 = B_2$.
  - **Taxman:** For a fixed $\tau \in [0, 1]$, we have $B'_1 = B_1 - b_1$ and $B'_2 = B_2 + (1 - \tau) \cdot b_1$.
- **All-pay:** Both players pay their bids.
  - **Richman:** $B'_1 = B_1 - b_1 + b_2$ and $B'_2 = B_2 + b_1 - b_2$. Thus, Player 1 pays Player 2 the difference between the two bids.
  - **Poorman:** $B'_1 = B_1 - b_1$ and $B'_2 = B_2 - b_2$.

Discrete vs. continuous bidding. A third orthogonal property of the bidding mechanism concerns the type of bids that are allowed: in *discrete bidding*, the budgets are given in coins and the minimal positive bid a player can make is one coin. In this survey, apart from Section 5, we consider *continuous bidding*, in which bids can be arbitrary small.

▶ Remark 1 (Ties in biddings). In continuous bidding, we avoid the issue of ties in the questions that we study (see Def. 3 for example). So while one needs to determine how ties are resolved, our results do not depend on a specific tie-breaking mechanism (apart from the corner cases in Thm. 32). Ties cannot be avoided in discrete bidding as we discuss in Section 5. ⌟

▶ **Definition 2** (Budget ratio). *For $i \in \{1, 2\}$, let B_i denote Player i's budget. Player i's ratio is then $B_i/(B_1 + B_2)$. In Richman bidding, since the bids only exchange hands, the sum of budgets is constant and we normalize it to 1, so that a player's ratio coincides with his actual budget.*

2 Qualitative First-Price Bidding Games

We consider the following qualitative objectives:

▶ **Definition 3** (Qualitative objectives). *Consider a graph with a set V of vertices.*

- ***Reachability:*** *There are two targets $t_1, t_2 \in V$. The game ends once a target t_i is visited, for $i \in \{1, 2\}$. Then, Player i is the winner.*
- ***Parity:*** *A parity index function is $p : V \to \mathbb{N}$. Player 1 wins an infinite play iff the maximal parity index that is visited infinitely often is odd.*

▶ Remark 4. Note that the definition of reachability objectives is slightly different than the typical definition in which only one player has a target and the second player (the *safety* player) wins iff the target is not reached. The main difference is that in the two-target version, there is no winner in a play in which both targets are not reached. As we show in Thm. 7 below, such ties do not occur. Thus, the two types of reachability objectives coincide in bidding games. See [4] for more details. ⌟

The main question studied in qualitative bidding games is the following:

What is a necessary and sufficient initial budget ratio that guarantees winning the game?

We illustrate a solution to this question in the following example.

▶ **Example 5.** Consider the reachability first-price Richman bidding game that is depicted in Figure 1. Player 1's goal is to reach t_1, and Player 2's goal is to reach t_2. We start with a naive solution by showing that Player 1 wins when his[1] budget exceeds 0.75. Suppose that the budgets are $\langle 0.75 + \varepsilon, 0.25 - \varepsilon \rangle$, for Player 1 and 2, respectively, for $\varepsilon > 0$. In the first turn, Player 1 bids 0.25 and wins the bidding since Player 2 cannot bid above $0.25 - \varepsilon$. The budgets are updated to $\langle 0.5 + \varepsilon, 0.5 - \varepsilon \rangle$ (in first-price Richman, only the higher bidder pays his bid to the other player). Player 1 moves the token to v_2. In the second bidding, Player 1 bids all his budget, wins the bidding since Player 2 cannot bid above $0.5 - \varepsilon$, moves the token to t_1, and wins the game.

A ratio of 0.75 thus suffices for winning, but is it necessary? No. The necessary and sufficient budget is in fact 2/3. More formally, we show that for every $\varepsilon > 0$, Player 1 can win when the initial budgets are $\langle 2/3 + \varepsilon, 1/3 - \varepsilon \rangle$. Player 1's first bid is 1/3, thus he wins the wins the bidding and moves the token to v_2. The budgets are updated to $\langle 1/3 + \varepsilon, 2/3 - \varepsilon \rangle$. Next, Player 1 bids all his budget, namely $1/3 + \varepsilon$. If he wins the bidding, he proceeds to t_1 and wins the game. Otherwise, Player 2 wins the bidding, and moves the token back to v_0. The key observation is that since Player 2 overbids, she bids at least $1/3 + \varepsilon$, thus Player 1's new budget is at least $2/3 + 2\varepsilon$. In other words, we are back to v_0 only that Player 1's increases by ε. By continuing in a similar manner, Player 1 forces his budget to increase by a constant every time we return to v_0. Thus, eventually, his budget exceed 0.75 and he can use the strategy above to force the game to t_1. A symmetric argument shows that Player 2 wins when the budgets are $\langle 2/3 - \varepsilon, 1/3 + \varepsilon \rangle$. ⌟

▶ **Definition 6** (Threshold ratios). *Consider a bidding game over a set of vertices V and let $B_1 \in [0, 1]$ be Player 1's initial budget ratio. The threshold ratio in a vertex $v \in V$, denoted $\mathit{Th}(v) \in [0, 1]$ is such that*

- *when $B_1 > \mathit{Th}(v)$, Player 1 can win the game from v, and*
- *when $B_1 < \mathit{Th}(v)$, Player 2 can win the game.*

[1] For fairness, we refer to Player 1/Max as *he* and Player 2/Min as *she*.

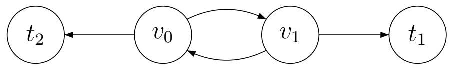

	t_2	v_0	v_1	t_1
Richman	1	2/3	1/3	0
poorman	1	$\frac{\sqrt{5}-1}{2}$	$\frac{3-\sqrt{5}}{2}$	0

Figure 1 A reachability bidding game with the threshold ratios under first-price Richman and poorman bidding.

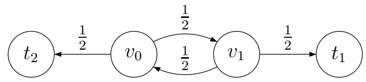

t_2	v_0	v_1	t_1
0	1/3	2/3	1

Figure 2 The random-turn game that corresponds the game in Fig. 1 with the probabilities to reach t_1 from each vertex.

2.1 Reachability first-price bidding games

The following theorem identifies threshold ratios in reachability first-price bidding games. See Fig. 1 for a specific instantiation of the theorem.

▶ **Theorem 7** ([23]). *Consider a reachability bidding game with target states t_1 and t_2. Threshold ratios exist under taxman bidding. Moreover, $Th(t_1) = 0$ and $Th(t_2) = 1$, and for all other vertices v, let v^- and v^+ be the neighbors of v such that $Th(v^-) \leq Th(v') \leq Th(v^+)$, for every neighbor v' of v. Then:*

- *under Richman bidding, $Th(v) = \frac{1}{2}(Th(v^+) + Th(v^-))$,*
- *under poorman bidding, $Th(v) = \frac{Th(v^+)}{1-Th(v^-)+Th(v^+)}$, and*
- *under taxman bidding with constant $\tau \in [0,1]$, we have $Th(v) = \frac{Th(v^+)+Th(v^-)-Th(v^-)\cdot\tau}{2-(1+Th(v^-)-Th(v^+))\cdot\tau}$.*

Proof. We sketch the proof for Richman bidding to give a flavor of the techniques. Suppose the initial vertex is v and that Player 1's budget is $B = \mathtt{Th}(v) + \varepsilon$. Player 1 maintains the invariant that whenever the token is placed on a vertex u, his ratio exceeds $\mathtt{Th}(u)$. The invariant holds initially. Suppose Player 1 bids $b = \frac{1}{2} \cdot \big(\mathtt{Th}(v^+) - \mathtt{Th}(v^-)\big)$, and upon winning, moves the token to v^-. If Player 1 wins the bidding, since $\mathtt{Th}(v) - b = \mathtt{Th}(v^-)$, the invariant is maintained at v^-. If he loses the bidding, Player 2 pays him at least b, and since $\mathtt{Th}(v) + b = \mathtt{Th}(v^+)$ and $\mathtt{Th}(v^+) \geq \mathtt{Th}(v')$, for every neighbor v', no matter where Player 2 moves, the invariant is maintained. The invariant implies that Player 1 does not lose. Indeed, recall that the sum of budgets is 1 in Richman bidding. Since $\mathtt{Th}(t_2) = 1$, the game cannot reach t_2, since then Player 1 owns more than the sum of budgets. To win the game, Player 1 plays as in Example 5: he uses his "excess" budget ε to bid slightly more than b so that he guarantees not losing while accumulating budget as the game proceeds. Eventually he accumulates enough budget to win $|V|$ biddings in a row and draws the game to t_1. ◀

Thm. 7 implies an intriguing equivalence between reachability Richman-bidding games and random processes. Before stating it in full generality, we illustrate the "plain vanilla" version of the equivalence, which applies to games in which all vertices have out-degree at most 2.

▶ **Example 8.** Consider the game depicted in Fig. 1. We construct a Markov chain by labeling all edges with probability 0.5 (see Fig. 2). Consider the probability of reaching the target t_1 from a vertex u, denoted $\mathbb{P}[\text{reach}(u, t_1)]$. Clearly, we have $\mathbb{P}[\text{reach}(t_1, t_1)] = 1$ and $\mathbb{P}[\text{reach}(t_2, t_1)] = 0$. As for the other two vertices, we have $\mathbb{P}[\text{reach}(v_0, t_1)] = \frac{1}{2}\mathbb{P}[\text{reach}(v_1, t_1)] + \frac{1}{2}\mathbb{P}[\text{reach}(t_2, t_1)]$, which is technically the same as the expression in Thm. 7 for threshold ratios in Richman bidding. Since the values for the targets are reversed, for every vertex u, we have $\mathtt{Th}(u) = 1 - \mathbb{P}[\text{reach}(u, t_1)]$. ⌟

Graphs with out-degree 2 are easier to reason about since it is clear which vertices are v^- and v^+ (which vertex is which does not matter, since we take a simple average between the two). In general, however, determining which are the two "important" vertices out of all the neighbors is not trivial.

▶ **Definition 9** (Random-turn game). *Consider a bidding game $\mathcal{G}$ that is played on a graph over a set of vertices V. For $p \in [0, 1]$, the random-turn game that corresponds to $\mathcal{G}$ w.r.t. p, denoted $\mathit{RT}(\mathcal{G}, p)$, is a game in which instead of bidding, in each turn we toss a (biased) coin to determine which player gets to move the token: Player 1 is chosen with probability p and Player 2 with probability $1 - p$. Formally, $\mathit{RT}(\mathcal{G}, p)$ is a stochastic game [17]. Every vertex $v \in V$, is replaced by three vertices v_N, v_1, and v_2. The vertex v_N simulates the coin toss, thus it has two outgoing edges: one with probability p to v_1 and a second with probability $1-p$ to v_2. For $i \in \{1, 2\}$, the vertex v_i is controlled by Player i and there are deterministic edges from v_i to u_N, for every neighbor u of v. The objective in $\mathit{RT}(\mathcal{G}, p)$ matches that of $\mathcal{G}$. For example, when $\mathcal{G}$ is a reachability bidding game, then $\mathit{RT}(\mathcal{G}, p)$ is a simple stochastic game.*

A *positional strategy* f in a stochastic game is a function that maps each vertex to a successor. When Player i, for $i \in \{1, 2\}$, plays according to f and the game reaches a vertex v that he controls, he moves the token to $f(v)$. Fixing two positional strategies in a stochastic game gives rise to a Markov chain. We restrict attention to positional strategies since in stochastic games with the objectives that we consider, existence of optimal positional strategies is well known.

▶ **Definition 10** (Values in simple stochastic games). *Consider a simple stochastic game $\mathcal{H}$ and a vertex v in $\mathcal{H}$. The* value *of v in $\mathcal{H}$, denoted $val(\mathcal{H}, v)$, is the probability that Player 1 reaches his target under optimal play of the two players.*

We state the general equivalence between the two models in the following theorem.

▶ **Theorem 11** ([24]). *Consider a reachability first-price Richman bidding game $\mathcal{G}$ over the vertices V. For every vertex $v \in V$, we have $\mathit{Th}(v) = 1 - val\big(\mathit{RT}(\mathcal{G}, 0.5), v\big)$.*

▶ Remark 12. We note that in graphs of finite size, threshold ratios under Richman bidding are always rational numbers; indeed, they coincide with values of a stochastic game, which in turn are a solution to a linear program. Since for poorman bidding the thresholds can be irrational (see Fig. 1), it seems unlikely that an equivalence with random-turn games exists. ⌟

2.2 Parity first-price bidding games

The solution to parity bidding games is based on the following lemma.

▶ **Lemma 13** ([7]). *Consider a reachability first-price taxman-bidding game $\mathcal{G}$ over the vertices V such that from every vertex, only Player 1's objective is reachable. Then, Player 1 wins with any positive initial budget from every vertex $v \in V$; thus, $\mathit{Th}(v) = 0$.*

We proceed to solve parity games.

▶ **Theorem 14** ([7]). *Parity first-price taxman-bidding games are linearly reducible to reachability taxman-bidding games. In particular, threshold ratios exists.*

Proof. Consider a bidding game $\mathcal{G}$. we identify the bottom strongly-connected components (BSCCs) of $\mathcal{G}$. We claim that in each such BSCC S there is a unique winner. The claim implies the theorem since the threshold ratios in the vertices of S are either all 0 or all 1. We then construct a reachability game on the rest of the graph in which each player's goal is to draw the game to a BSCC that is winning for him.

To prove the claim, suppose the maximal parity index in a BSCC S is obtained in a vertex t and that it is even, and the other case is dual. We think of S as a reachability game in which Player 2's target is t and Player 1 has no target. By Lem. 13, Player 2 can force a visit to t with any positive initial budget. To force infinite many visits, she splits her initial budget $\varepsilon > 0$ into infinite many parts $\varepsilon_1, \varepsilon_2, \ldots$, and uses a budget of ε_i, for $i \geq 1$, to force a visit to t for the i-th time. ◀

2.3 Computational complexity and open problems

Intuitively, the computational problem that we would like to solve is finding the threshold ratio in a vertex. Formally, one way to define the corresponding decision problem is given a game $\mathcal{G}$ and a vertex v in $\mathcal{G}$, decide whether $\mathtt{Th}(v) > 0.5$.

▶ **Theorem 15** ([23, 4, 5, 7])**.** *For parity Richman bidding, finding threshold ratios is in NP ∩ coNP, and it is in P when all vertices have outdegree at most 2 or when the graph is undirected. For taxman and poorman bidding the problem is in PSPACE.*

Proof (Sketch). The upper bound for Richman bidding immediately follows from the equivalence with random-turn games in Thm. 11. Indeed, random-turn games are a special case of stochastic games, and solving the later is in NP ∩ coNP [17]. When the outdegree is 2, the solution is obtained by solving a linear program. An algorithm for undirected graphs was shown in [23] (and was extended recently to solve biased undirected random-turn games [28]). For taxman bidding, the properties of threshold ratios in Thm. 7 are polynomial but not linear, thus we reduce the problem to the *existential theory of the reals* [15]. ◀

▶ **Open problem 1** (Tightening the computational complexity gap)**.** *There is a gap in our understanding in terms of computational complexity. It is possible that solving Richman bidding games is in P, and it is possible that they are as hard as solving stochastic games. Possible lower bounds for poorman and taxman bidding include showing that the problem is NP-hard or as hard as solving the existential theory of the reals.*

3 Mean-Payoff First-Price Bidding Games

Before formalizing the mean-payoff objective, we motivate it in the following application.

▶ **Example 16.** Suppose an internet content-provider (e.g., New York Times) has one ad slot for sale on its website. Two advertisers repeatedly (e.g., daily) compete to publish their ad in a first-price auction that the content provider holds. Our goal is to find an optimal bidding strategy for an advertiser that, given his budget constraints, maximizes the long-run ratio of the time that his ad shows. To find such a strategy, we reason about the Bowtie game in Fig. 3; we associate our advertiser with Player 1. Whenever Player 1 wins a bidding, he moves to v_{Max}, which represents his ad being shown that day. Informally, Player 1's goal is to maximize the time his ad shows (e.g., maximize the number of days his ad appears in a year). ⌟

Formally, a *mean-payoff game* is played on a weighted graph $\langle V, E, w\rangle$, where $w : V \to \mathbb{Q}$. Each infinite play has a *payoff*, which is Player 1's reward and Player 2's cost, thus we call the players in a mean-payoff game Max and Min, respectively.

▶ **Definition 17** (Payoff and energy). *Consider an infinite path $\eta = \eta_0, \eta_1, \ldots$. For $n > 1$, let $\eta^n = \eta_0, \ldots, \eta_n$ be a prefix of η. The* energy *of η^n, denoted $\mathsf{energy}(\eta^n)$, is the sum of weights it traverses, thus $\mathsf{energy}(\eta^n) = \sum_{0\leq i<n} w(\eta_i)$. The* payoff *of η, denoted $\mathsf{payoff}(\eta)$, is $\mathsf{payoff}(\eta) = \liminf_{n\to\infty} \frac{1}{n} \cdot \mathsf{energy}(\eta^n)$. Note that the use of $\liminf$ gives Min an advantage.*

We consider the following main question when studying mean-payoff bidding games.

Given an initial budget ratio r, what is the optimal payoff a player can guarantee?
For example, consider the Bowtie game in Fig. 3. What is the optimal payoff Max can guarantee with a budget ratio of $2/3$ under Richman bidding? How does it compare with the optimal payoff under poorman bidding? Does the answer change when the ratio is $1/3$?

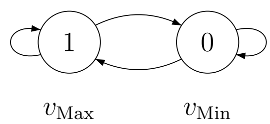

Figure 3 The Bowtie mean-payoff game $\mathcal{G}_{\bowtie}$.

We answer the questions in full generality after the following definitions. The specific solutions in the Bowtie game can be found in Example 21.

▶ **Definition 18** (Mean-payoff value in bidding games). *Consider a strongly-connected mean-payoff bidding game $\mathcal{G}$ and a budget ratio $r \in (0, 1)$. The* mean-payoff value *of $\mathcal{G}$ w.r.t. r, denoted $MP(\mathcal{G}, r)$, is $c \in \mathbb{R}$ if independent of the initial vertex,*

- *when Max's initial ratio exceeds r, he[1] has a strategy that guarantees a payoff of $c - \varepsilon$, for every $\varepsilon > 0$, and*
- *Max cannot do better: with a ratio that exceeds $1 - r$, Min can guarantee a payoff of at most $c + \varepsilon$, for every $\varepsilon > 0$.*

▶ **Definition 19** (Mean-payoff value in random-turn games). *For $p \in [0, 1]$, since $\mathcal{G}$ is a mean-payoff game, the random-turn game $RT(\mathcal{G}, p)$ is a stochastic mean-payoff game. The expected payoff under optimal play of the two players is called the mean-payoff value, which we denote $MP(RT(\mathcal{G}, p))$. The value is known to exist [30], and since $\mathcal{G}$ is strongly-connected, the value does not depend on the initial vertex.*

The theorem below states the equivalence between mean-payoff first-price bidding games and random-turn games.

▶ **Theorem 20** ([5, 4, 7]). *Consider a strongly-connected mean-payoff game $\mathcal{G}$ and a ratio $r \in (0, 1)$. Then:*

- *under Richman bidding, $MP(\mathcal{G}, r) = MP\big(RT(\mathcal{G}, 0.5)\big)$,*
- *under poorman bidding, $MP(\mathcal{G}, r) = MP\big(RT(\mathcal{G}, r)\big)$, and*
- *under taxman bidding with constant $\tau \in [0, 1]$, we have $MP(\mathcal{G}, r) = MP\big(RT(\mathcal{G}, \frac{r+\tau\cdot(1-r)}{1+\tau})\big)$.*

▶ **Example 21.** We apply Thm. 20 to the bowtie mean-payoff game $\mathcal{G}_{\bowtie}$ (Fig. 3). Since the outdegree in $\mathcal{G}_{\bowtie}$ is 2, the random-turn game that corresponds to it is simple and depicted in Fig. 4. Informally, we expect a random run to "stay" in v_{Max} portion p of the time, and since the weights are simple, we have $\mathsf{MP}(\mathsf{RT}(\mathcal{G}_{\text{Bowtie}}, p)) = p$.

Under Richman bidding, the initial ratio does not matter, the equivalence is always with the fair random-turn game, thus $\mathsf{MP}(\mathcal{G}_{\bowtie}, r) = 0.5$, for every $r \in (0, 1)$. We find the result for mean-payoff poorman bidding more surprising, since no equivalence is known for reachability poorman bidding games (see Remark 12). With an initial ratio of $r \in (0, 1)$, the optimal payoff under poorman bidding is r. Thus, Max prefers poorman over Richman when his ratio is 2/3 and prefers Richman over poorman when the ratio is 1/3. Interestingly, when the ratio is 1/2, the payoffs under poorman and Richman bidding coincide. ⌟

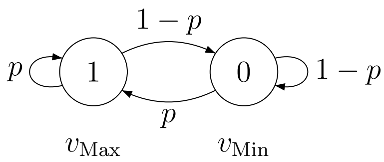

Figure 4 The random-turn game that corresponds to $\mathcal{G}_{\bowtie}$ in Fig. 3 w.r.t. $p \in [0, 1]$.

▶ Remark 22 (Strategies in bidding games vs. in random-turn games). We point out that strategies in bidding games are much more complicated than in stochastic games. At a vertex v in a stochastic game, a strategy only needs to select a vertex u to move the token to from v. In a bidding game, in addition to the choice of u, a strategy prescribes a bid b. As the proof of Thm. 7 shows, in reachability games, knowing the threshold ratios and assuming they are positive, immediately gives us both u and b. In mean-payoff games, this is no longer the case. Indeed, the statement of Thm. 20 only tells us the optimal payoff a player can guarantee. Knowing the optimal payoff does not hint at how to achieve it. The proof of Lem. 23 demonstrates that finding the right bids is not a trivial task. In fact, there is an alternative existential proof for the Richman part of Thm. 20 (see [4]). The proof relies on the equivalence shown in [23] between reachability Richman bidding games and random-turn games adapted to infinite graphs together with results on *one-counter simple-stochastic games* [14, 13]. Beyond the lack of bidding strategy, since no equivalence is known for reachability objectives apart from Richman bidding, the existential proof does not extend to poorman or taxman bidding. ⌟

3.1 Solving the Bowtie game

In the next two sections, we give a flavor of the techniques used to prove Thm. 20. In this section, we solve $\mathcal{G}_{\bowtie}$ and in the next section, we describe a framework to extend a solution to $\mathcal{G}_{\bowtie}$ to general SCCs. See Open problem 3 for a discussion on the asymmetry between Min and Max in the following lemma.

▶ **Lemma 23.** *Consider the mean-payoff game $\mathcal{G}_{\bowtie}$ (Fig. 3) under Richman bidding. Irrespective of the initial ratios:*
- *Min has a strategy that guarantees a payoff of at most* 0.5.
- *For every $\varepsilon > 0$, Max has a strategy that guarantees a payoff of at least* $0.5 - \varepsilon$.

Proof. An optimal Min strategy. The details of the construction can be found in [4]. It is convenient to re-normalize the weights and set $w(v_{\text{Min}}) = -1$ so that $\mathsf{MP}\big(\mathsf{RT}(\mathcal{G}_{\bowtie}, 0.5)\big) = 0$. Recall that the energy of a finite path, which we denote with k below, is the sum of weights that it traverses. The construction is based on the following observation whose proof is a simple exercise.

Observation: Suppose Min plays according to a strategy that guarantees that in every infinite play either the energy (1) equals 0 infinitely often, or (2) is bounded from above by a constant. Then, the payoff is non-negative.

Suppose that the energy starts from $k_0 > 0$ and Min's initial budget is $B^I_{\text{Min}} > 0$. Min chooses $N \in \mathbb{N}$ such that $B^I_{\text{Min}} = k_0/N + \delta$, for some $\delta > 0$. This is possible since B^I_{Min} and k_0 are constants. We think of δ as "spare change" that Min does not use, but will be helpful at the end of the proof. We call the rest of her budget the *main* budget.

Min's strategy maintains the following invariant: when the energy of a finite play is k, Min's budget exceeds $\frac{k}{N}$. The invariant holds initially by our choice of N. Suppose that Min's main budget following a finite play with energy k is $\frac{k}{N}$. Min bids $\frac{1}{N}$. We distinguish between two cases. If Min wins the bidding, the energy drops to $k-1$ and since she pays $\frac{1}{N}$, her main budget drops to $\frac{k-1}{N}$. On the other hand, if she loses, the energy increases to $k+1$, and Max pays her at least $\frac{1}{N}$, thus her main budget increases to $\frac{k+1}{N}$. The strict inequality in the invariant is obtained from the spare change δ. Finally, recall that in Richman bidding the sum of budgets is 1. Thus, the invariant ensures that the energy is bounded from above by N, since by plugging $k = N$ in the invariant, Min's budget would exceed 1.

To conclude the construction, Min plays as follows. When the energy is 0, Min bids 0 and "waits" for the energy to increase. When it increases to k_0^1, she plays according to the strategy above to guarantee a bounded energy N_1. When the energy drops to 0, Min "waits" until the energy increases again to k_0^2, and she plays to keep the energy bounded by N_2, and so on. If there is an $n \geq 1$, such that the energy hits 0 only $n-1$ times, then the energy is bounded from above by N_n. Otherwise, the energy hits 0 infinitely often. Thus, by the observation, the payoff is at most 0.

An optimal Max strategy. An explicit construction for Max was shown in [4] and an existential one in [7]. We describe a simpler construction, which is due to Ismaël Jecker [9].

Let $\varepsilon > 0$. This time, we re-normalize the weights to be $w(v_{\text{Max}}) = c = 1+\varepsilon$ and $w(v_{\text{Min}}) = -1$.

Observation: Suppose Max plays according to a strategy that guarantees that the energy is bounded from below by a constant. Then, the payoff with the updated weights is non-negative, and the payoff with the original weights is at least $\frac{1}{2+\varepsilon}$.

Let α such that $(1+\alpha) = (1-\alpha)^{-c}$, and such α exists since $c > 1$. Max maintains the invariant that when the energy is $k \in \mathbb{N}$, his budget exceeds $(1+\alpha)^{-k}$. The invariant implies $k > 0$, since otherwise, Max's budget is greater than $(1+\alpha)^0 = 1$, which, again, is impossible since the sum of budgets in Richman bidding is 1. We choose an initial energy level $k_0 \in \mathbb{N}$, thus the energy will in fact be bounded from below by $-k_0$. We choose k_0 such that $B^I_{\text{Max}} = B + \delta$, where $B = (1+\alpha)^{-k_0}$ and $\delta > 0$. This is possible since $\lim_{k\to\infty}(1+\alpha)^{-k} = 0$. Again, the "spare change" δ is never used and we refer to B as Max's main budget. Max maintains the invariant that when the energy is k, his main budget is at least $(1+\alpha)^{-k}$. When Max's main budget is B, he bids $\alpha \cdot B$. We distinguish between the two outcomes of a bidding. If Max loses, the energy decreases to $k-1$. Moreover, Min overbids Max, thus Max's new main budget B' is at least $B + \alpha B \geq \frac{1}{(1+\alpha)^k} + \frac{\alpha}{(1+\alpha)^k} = (1+\alpha)^{k-1}$. On the other hand, if Max wins, the energy increases to $k+c$ and his new main budget B' is at least $B - \alpha B = \frac{1-\alpha}{(1+\alpha)^k}$. Since $(1+\alpha) = (1-\alpha)^{-c}$, we obtain $(1-\alpha) = (1+\alpha)^{-c}$, thus $B' = (1+\alpha)^{-(k+c)}$, as required. ◀

Poorman bidding

An explicit construction of an optimal strategy for Max under poorman bidding was shown in [5], an existential construction was shown in [7], and a significantly simpler explicit construction was shown in [9]. Those constructions use similar ideas to the construction shown in Lem. 23 for Max, though they are technically more involved. We leave the constructions out of this survey for sake of brevity. We would still like to convince the reader that contrary to Richman bidding, the initial ratio matters in poorman bidding; i.e., with a higher ratio a player can guarantee a better payoff. Below we show a simple strategy construction that is optimal for Min in the Bowtie game.

▶ **Lemma 24** ([5]). *Suppose that Min's initial budget is $\ell \in \mathbb{N}$ and Max's initial budget is 1. Then, in $\mathcal{G}_{\bowtie}$, Min can guarantee a payoff of at most $\frac{\ell}{\ell+1}$.*

Proof. We re-normalize the weights so that $w(v_{\text{Min}}) = -1$ and $w(v_{\text{Max}}) = 1$. Suppose first that both budgets are 1. We describe the "tit for tat" strategy for Min that, intuitively, copies Max's bidding strategy. We assume for simplicity that Min wins bidding ties. Formally, Min's strategy is based on a queue of numbers as follows: if the queue is empty, Min bids 0, and otherwise she bids the maximal number in the queue. If Min wins with a bid of $b > 0$, she removes b from the queue. If Max wins with a bid of $b > 0$, Min adds b to the queue.

We make several observations. (1) Min's strategy is legal: it never bids higher than the available budget; indeed, Max bids legally, the budgets are the same, and Min's sum of winning bids is at most Max's sum of winning bids. (2) The size of the queue is an upper bound on the energy; indeed, every bid in the queue corresponds to a Max winning bid that is not "matched" (the size is an upper bound since Min might win biddings when the queue is empty). (3) If Min's queue fills, it will eventually empty; indeed, if $b \in \mathbb{R}$ is in the queue, in order to keep b in the queue, Max must bid at least b, thus eventually his budget runs out. Combining the three, since the energy is at most 0 when the queue empties, Min's strategy guarantees that the energy is at most 0 infinitely often and as in Lem. 23, the payoff is non-positive.

Now, assume Min's budget is $\ell \in \mathbb{N}$. Min's strategy is the same only that whenever Max wins with $b > 0$, Min adds ℓ copies of b to the queue. The strategy is legal since, intuitively, we can think of Min's budget as if it consists of ℓ parts of size 1, and each copy of the bid b is paid out of a unique part. Legality then follows as in the simple case above. The other two observations above still hold, only that now, since every Max win is matched by ℓ Min wins, when the queue empties, the number of Min wins is at least ℓ times as much as Max's wins. It follows that the payoff is at most $\frac{\ell}{\ell+1}$. ◀

3.2 Solving strongly-connected games

In this section, we describe a general framework that, intuitively, allows us to extend a solution for $\mathcal{G}_{\bowtie}$ to any strongly-connected mean-payoff game. We describe the challenges of the extension in the following example.

▶ **Example 25.** It is helpful to think about changes in the budget and energy throughout a play as two "walks" on two sequences; a "budget sequence" and an "energy sequence". We preserve an invariant between the two walks. Recall, for example, that in Lem. 23, when the energy is k, Max's ratio is $(1+\alpha)^{-k}$. When Max wins a bidding, the energy walk takes a step to $k+1$ and the budget walk takes a step down to $(1+\alpha)^{-(k+1)}$. The key idea that concludes the proof is that the budget walk is bounded from above, e.g., in Richman bidding the sum of budgets is 1. Due to the invariant, the bound on the budget implies a bound on the energy and in turn a guarantee on the payoff.

The structure of $\mathcal{G}_{\bowtie}$ is convenient since in each bidding, both walks always take exactly one step at a time. In general, this is not the case. For example, consider the game depicted in Fig. 5. Cycling v_4 decreases the energy by 2, so to even out, the budget walk should also take two steps at once. Moreover, there are more involved cycles. Suppose the game starts from v_3. Note that Max always moves left (see the red edges) and Min moves to the right. How should Max bid in v_3 to maintain the invariant between budget and energy? On the one hand, if he wins the first bidding and loses the second, the cycle v_3, v_2, v_3 forms with a change in energy of -2. On the other hand, if he loses the first bidding and wins the second, the cycle $v_3 v_4 v_3$ forms with a change of energy of -3. Finally, there is no guarantee that the game returns to v_3 ⌟

To overcome these challenges, we introduce a measure of "importance" to vertices, which we call the *strength*. We need several definitions. Let $p \in [0,1]$. Inspired by the notation in reachability games, we denote by v^- and v^+ the optimal vertices to move to from v for Min and Max, respectively. Formally, for optimal memoryless strategies σ_{Max} and σ_{Min} in $\mathsf{RT}(\mathcal{G}, p)$, we define $v^- = \sigma_{\text{Min}}(v)$ and $v^+ = \sigma_{\text{Max}}(v)$. The concept of *potentials* was originally defined in the context of the strategy iteration algorithm [20]. 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 solutions to the following equations:

$$\text{Pot}_p(v) = p \cdot \text{Pot}_p(v^+) + (1-p) \cdot \text{Pot}_p(v^-) + w(v) - \mathsf{MP}(\mathsf{RT}(\mathcal{G}, p))$$
$$\text{St}_p(v) = p \cdot (1-p) \cdot \big(\text{Pot}_p(v^+) - \text{Pot}_p(v^-)\big)$$

We suggest an intuitive way to read the definition of potentials. Consider a weighted Markov chain in which each vertex v has two neighbors v^+ and v^-, and the probability of proceeding from v to v^+ and v^- is respectively p and $1-p$. Suppose there is a target that is reached with probability 1. Let $E(v)$ denote the expected energy in a path from v to the target. We note that $E(v)$ roughly coincides with $\text{Pot}_p(v)$. Indeed, to compute $E(v)$, we take the current energy, i.e., $w(v)$, and add, for each neighbor u, the probability of proceeding from v to u multiplied by the expected energy to the target from u, thus we obtain $E(v) = w(v) + p \cdot E(v^+) + (1-p) \cdot E(v^-)$.

The following lemma connects potential, energy, and strength and thus gives rise to the invariant between budget and energy. We need several definitions. Consider a finite path $\eta = v_1, \ldots, v_n$ in $\mathcal{G}$. We intuitively think of η as a play, where for every $1 \le i < n$, the bid of Max in v_i is $\text{St}_p(v_i)$ and he moves to v_i^+ upon winning. Thus, when $v_{i+1} = v_i^+$, we think of Max as *investing* $\text{St}_p(v_i)$ and when $v_{i+1} \neq v_i^+$, we think of Min winning the bid thus Max *gains* $\text{St}_p(v_i)$. We denote by $I(\eta)$ and $G(\eta)$ the sum of investments and gains, respectively. Recall that the energy of η is $\mathsf{energy}(\eta) = \sum_{0 \le i < n} w(v_i)$.

▶ **Lemma 26** ([4, 5, 7]). *Consider a strongly-connected game $\mathcal{G}$, and $p = \frac{\nu}{\nu+1} \in (0,1)$, thus Max's budget is ν and Min's budget is 1. For every finite path $\eta = v_1, \ldots, v_n$ in $\mathcal{G}$, we have*

$$\mathit{Pot}_p(v_1) - \mathit{Pot}_p(v_n) + (n-1) \cdot \mathit{MP}(\mathit{RT}(\mathcal{G}, p)) \le \mathit{energy}(\eta) + G(\eta) \cdot \nu - I(\eta).$$

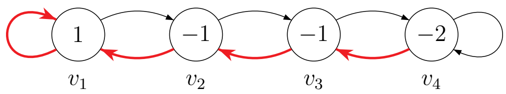

	v_1	v_2	v_3	v_4
Potential:	6	3	0	−3
Strength:	1	2	2	1

Figure 5 A mean-payoff game $\mathcal{G}$ with $\mathsf{MP}(\mathsf{RT}(\mathcal{G}, 2/3)) = 0$, Max's optimal moves are depicted in red, and $\text{Pot}_{2/3}$ and $\text{St}_{2/3}$ below the vertices.

▶ **Example 27.** Consider the game depicted in Fig. 5. Max always proceeds left, that is, for every vertex v, the vertex v^+ is the one to its left. The potentials and the resulting strengths are depicted below the vertices. It is not hard to verify that the stationary distribution of $\mathsf{RT}(\mathcal{G}, 2/3)$ is $\langle \frac{8}{15}, \frac{4}{15}, \frac{2}{15}, \frac{1}{15} \rangle$. Multiplying by the weights, we obtain $\mathsf{MP}(\mathsf{RT}(\mathcal{G}, 2/3)) = 0$.

It is clearest to exemplify Lem. 26 in cycles, where the left-hand side of the inequality completely cancels out. Consider the cycle $\eta = v_3, v_2, v_1, v_1, v_2, v_3$ in which Max wins the first three bids, thus $I(\eta) = 2 + 2 + 1 = 5$, and loses the last two biddings, thus $G(\eta) = 1 + 2$. We have $\mathsf{energy}(\eta) = -1$ since the last vertex on the path does not contribute to the energy. As Lem. 26 predicts, we have $I(\eta) - G(\eta) \cdot \nu = 5 - 3 \cdot 2 = -1 = \mathsf{energy}(\eta)$. We encourage the reader to verify that every cycle has this property. ⌟

3.3 Solving general games, computational complexity, and open problems

As in the case of qualitative objectives, there is a gap in our understanding of the computational complexity of mean-payoff games. We state the known upper bounds.

▶ **Theorem 28** ([4, 5, 7]). *Given a mean-payoff game $\mathcal{G}$, a vertex v in $\mathcal{G}$, a ratio $r \in (0, 1)$, and a constant $c \in \mathbb{Q}$, deciding whether Max can guarantee a payoff greater than c from v with an initial ratio that exceeds r is:*
- *In PSPACE for taxman bidding and general graphs.*
- *In NP $\cap$ coNP for Richman bidding, or with taxman bidding when $\mathcal{G}$ is strongly-connected.*
- *In P when $\mathcal{G}$ is strongly-connected and the out-degree of all vertices is 2.*

Proof (Sketch). Thm. 20 allows us to reduce general mean-payoff games to reachability games. Under Richman bidding, the theorem implies that every bottom strongly-connected component (BSCC) $\mathcal{S}$ of $\mathcal{G}$ has a "winner": we say that $\mathcal{S}$ is winning for Max iff $\mathsf{MP}\big(\mathsf{RT}(\mathcal{S}, 0.5)\big) > c$. Indeed, all Max needs to do, is to reach $\mathcal{S}$ with a positive ratio. Upon reaching $\mathcal{S}$, he can switch to a strategy that guarantees $\mathsf{MP}\big(\mathsf{RT}(\mathcal{S}, 0.5)\big)$. Thus, similarly to Thm. 14, we construct a reachability game in which the targets for Max are his winning BSCCs and the targets for Min are her winning BSCCs. Membership in NP and coNP for Richman bidding is obtained by guessing the winning BSCCs and guessing positional strategies both in the reachability random-turn game and the mean-payoff random-turn games in the BSCCs. One can verify in polynomial time that the strategies are optimal and calculate the mean-payoff values that they achieve. Verifying that a BSCC is winning is then trivial.

For poorman and taxman bidding, the ideas are similar though more complicated. For a BSCC $\mathcal{S}$, we need to find a ratio $r(\mathcal{S})$ such that Max can guarantee a payoff greater than c with $r(\mathcal{S})$. The reachability game we construct is more involved: the game ends once a target $\mathcal{S}$ is reached, and Max wins iff his ratio is greater than $r(\mathcal{S})$. Recall that in reachability bidding games, the threshold ratios of the targets is $\mathtt{Th}(t_1) = 0$ and $\mathtt{Th}(t_2) = 1$, and every other vertex it is some "average" of two of its neighbors that depends on the bidding mechanism. The solution here in the intermediate vertices is the same average as in reachability bidding games with the corresponding bidding mechanism. The only change is in the targets; namely, the threshold ratio in a target $\mathcal{S}$ is $r(\mathcal{S})$.

We turn to the last point. Under Richman bidding, since the random-turn game is un-biased, for every vertex v with two neighbors v_1 and v_2, we can construct a linear program without finding which of the vertices is v^+ and which is v^-. Indeed, the linear program ensures that $\mathtt{Th}(v) = \frac{1}{2} \cdot (\mathtt{Th}(v_1) + \mathtt{Th}(v_2))$. A similar solution applies to mean-payoff Richman bidding games. On the other hand, for poorman and taxman bidding, since the random-turn game is biased, we need to find which of the vertices is v^+ and which is v^-. Fortunately, it is shown in [5] that when the out-degree is 2, it suffices to solve an MDP rather than a stochastic game, hence the polynomial-time upper bound. ◀

Open problems

The bounds in Thm. 28 are not tight, thus Open problem 1 applies here as well. In addition, we state the following open problem.

▶ **Open problem 2** (The target-payoff problem). *Study the dual of the problem that is studied in Thm. 28: given a target payoff, find the minimal ratio that is required to guarantee it.*

▶ **Open problem 3** (ε-optimal vs. optimal strategies). *Consider a strongly-connected mean-payoff game $\mathcal{G}$ and an initial ratio r for Max. For all but Richman bidding, to prove Thm. 20, we construct* ε-optimal strategies *for Max, namely strategies that guarantee a payoff of $\mathsf{MP}(\mathcal{G}, r) - \varepsilon$, for every $\varepsilon > 0$. This is sufficient since the definition of mean-payoff (Def. 17) gives Min an advantage, thus an ε-optimal strategy for Min is obtained by relating Min with Max in the "dual" game in which the weights are multiplied by -1. We call a strategy* optimal *if it guarantees a payoff of $\mathsf{MP}(\mathcal{G}, r)$. This raises two open questions:*

- *Find optimal strategies for Max.*
- *Find optimal strategies for Min for poorman and taxman bidding.*[2]

Energy games. An *energy game* is played on a weighted directed graph. Min wins iff there is a prefix in which the energy hits 0. Energy games are qualitative games and so the main question regards the threshold ratio: given an initial vertex v and an initial energy level k, find the threshold ratio $\mathtt{Th}(v, k)$ that is necessary and sufficient for Min to win. Our solution to mean-payoff games has some implications on energy games. Consider a strongly-connected energy game $\mathcal{G}$. It is not hard to adapt Min's strategy in Lem. 23 to show that under Richman bidding, when $\mathsf{MP}(\mathsf{RT}(\mathcal{G}, 0.5)) = 0$, Min wins from every initial energy level and with any positive initial budget, thus $\mathtt{Th}(v, k) = 0$, for every $v \in V$ and $k \in \mathbb{N}$. On the other hand, if $\mathsf{MP}(\mathsf{RT}(\mathcal{G}, 0.5)) > 0$, then $\lim_{k\to\infty} \mathtt{Th}(v, k) = 1$, thus for every initial ratio $r > 0$, there is an initial energy level from which Max wins. A similar corollary holds for poorman and taxman bidding. In [23], the threshold ratios for the Bowtie game are classified under poorman bidding: for every $k \in \mathbb{N}$, we have $\mathtt{Th}(v, k) = \frac{k+2}{2k+2}$.

▶ **Open problem 4** (Energy bidding games). *Find solutions to other energy bidding games. For example, solving the Bowtie game with weights $+2$ and -1 under poorman bidding is open. The only thing we know is that when Max's budget exceeds $2/3$, there is an initial energy level from which he wins.*

4 All-Pay Bidding Games

Before describing the theoretical properties of all-pay bidding games, we illustrate applications of the model. Throughout the survey we always think of the players' budgets as money. Applications arise, however, from thinking of the budgets as bounded resources with little or no inherent value that the players use in order to achieve some goal. *Colonel Blotto* games, which date back to [12] and have been extensively studied since, have been used to

[2] Note that Lem. 24 constructs an optimal strategy for the Bowtie game under poorman bidding. We do not know how to extend this strategy to general strongly-connected games.

model such settings; e.g., rent seeking [32], patent races (e.g., [10]), political campaigning or lobbying, sport competitions, or biological auctions. The terminology there is that two colonels own armies and compete in one-shot in n battlefields; the colonels need to distribute their armies across the battlefields, in each battlefield the outnumbering army wins, and possible goals include maximizing the number of battlefields won or the probability that at least $k < n$ battlefields are won. In our terminology, the colonels are performing n concurrent biddings. It was argued in [21] that many of the examples above are actually dynamic in nature. We thus argue that all-pay poorman bidding games, which are a dynamic version of Colonel Blotto games, capture these settings in a more precise way.

Another application of bidding games is reasoning about systems in which the scheduler accepts payment in exchange for priority. Blockchain technology is one such system, where the *miners* accept payment and have freedom to decide the order of blocks they process based on the proposed transaction fees. Transaction fees are not refundable, thus all-pay poorman bidding is the most appropriate modelling. Manipulating transaction fees is possible and can have dramatic consequences: such a manipulation was used to win a popular game on Ethereum called FOMO3d[3]. There is thus ample motivation for reasoning and verifying blockchain technology [16].

Rechability all-pay poorman games

Reachability all-pay poorman bidding games were studied in [8]. The bad news is that reachability games under all-pay bidding are significantly more complicated than under first-price bidding. In light of the technical difficulty, it is encouraging that simple yet powerful results are obtained on this model. We focus mostly on the following class of games.

▶ **Definition 29** (Race games). *In the race game $G(i, j)$, for $i, j \in \mathbb{N}$, Player 1 wins iff he wins i biddings before Player 2 wins j biddings.*

▶ **Example 30.** Suppose for example that two NBA teams play a "best of 7" tournament. The resources (the athletes' strengths) need to be invested throughout the tournament so as to maximize the probability that the tournament is won. Note that the resources have no inherent value, that invested resources are exhausted rather than transferred between the teams, and that the budgets need not be equal (one the teams can be fresher or have a deeper bench). Thus, in order to find an optimal strategy to select how much to invest in each game, we solve the race game $G(4, 4)$ under all-pay poorman bidding. ⌟

The simplest interesting reachability game is $G(2, 1)$, in which Player 1 needs to win two biddings in a row. Throughout all of this section we fix Player 2's budget to 1. Under first-price bidding, the solution to this game is almost trivial: the threshold ratios are 2 and 3 under poorman and Richman bidding, respectively. Let B_1 be Player 1's budget. It was observed already in [23] that under all-pay poorman bidding, for $B_1 \in [1, 2]$, neither player has a deterministic winning strategy. We thus need the following definitions.

▶ **Definition 31.** *A* mixed *strategy assigns a probability distribution over bids in each step. The* lower value *in an all-pay game $\mathcal{G}$ w.r.t. an initial ratio r is $val^{\downarrow}(\mathcal{G}, r) = \sup_f \inf_g \int_{\pi \sim D(f,g)} \mathbb{P}[\pi$ is winning for Player 1$]$. The* upper value*, denoted $val^{\uparrow}(\mathcal{G}, r)$ is defined dually. It is always the case that $val^{\downarrow}(\mathcal{G}, r) \leq val^{\uparrow}(\mathcal{G}, r)$, and when $val^{\uparrow}(\mathcal{G}, r) = val^{\downarrow}(\mathcal{G}, r)$, we say that the value exists and we denote it by $val(\mathcal{G}, r)$.*

[3] `https://bit.ly/2wizwjj`

The following problem was left open in [23] already for $G(2,1)$:

Does the value exist in every vertex of every all-pay bidding game?
If it does, characterize the value.

We solve this problem for the game $G(2,1)$.

▶ **Theorem 32** ([8]). *Consider the all-pay bidding game $G(2,1)$ in which Player 1 needs to win twice in a row. The value exists for every pair of initial budgets, the optimal Player 1 strategy has finite support, and the value as a function of Player 1's budget is the following step-wise function: if Player 1's budget is B_1 and Player 2's budget is 1, then, when $B_1 > 2$, the value is 1, when $B_1 < 1$, the value is 0, and when $B_1 \in (1+\frac{1}{n+1}, 1+\frac{1}{n}]$, for $n \in \mathbb{N}$, then the value is $\frac{1}{n+1}$.*

Thm. 32 is positive in several aspects, however all-pay bidding games are not that simple.

▶ **Theorem 33** ([8]). *Consider the game $G(3,1)$ in which Player 1 needs to win three biddings in a row. Suppose the initial budgets are 1.25 for Player 1 and 1 for Player 2. Then, an optimal strategy for Player 1 requires infinite support in the first bidding.*

We turn to present simple yet powerful positive news about all-pay bidding games.

An approximation algorithm. We obtain an approximation algorithm for the upper- and lower-values in games played on DAGs. The algorithm is based on a discretization of the budgets similar in spirit to discrete bidding. That is, we restrict the granularity of the bids and require Player 1 to bid multiples of some $\varepsilon > 0$. The smaller the ε, the better the discrete game approximates the continuous game (see [8] for the formal guarantees). A discrete bidding game is a finite game (the smaller ε is, the larger the game). Finding values can thus be done in polynomial time using linear programming.

We implemented the algorithm and ran it on several race games, see the plots in Fig. 6. Each game is associated with two plots; a lower- and an upper-bound. It is encouraging that the difference between the two plots is barely visible already for $\varepsilon = 0.01$. That is, the plots hint that the lower- and upper-values coincide and that the value exists in the games that we experiment with. It is also reassuring to find an experimental confirmation for Thm. 32 (see the red plots). While the statement in Thm. 33 cannot be confirmed by the plots, the value function in the budget region of 1.25 seems involved (see the purple plots). For more complicated games, the plots hint that the value as a function of the budgets is continuous rather than step-wise as in $G(2,1)$.

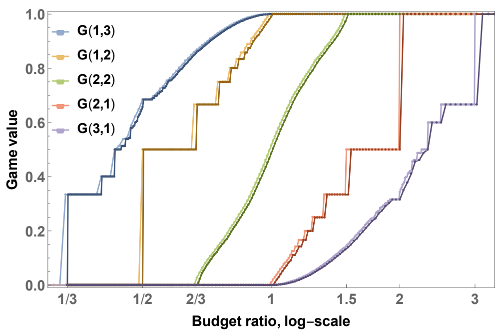

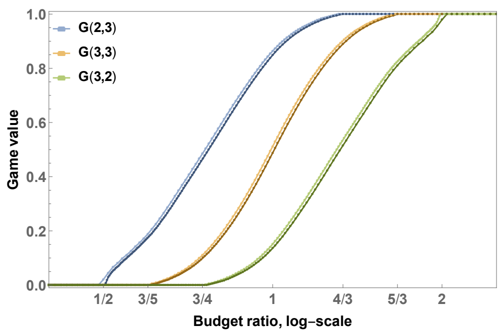

Figure 6 Approximations for the upper- and lower-values for several race games.

Surely-winning threshold ratios. we considering the question of *surely winning* in all-pay bidding games, namely winning with probability 1, and identify a threshold phenomena similar to first-price bidding.

▶ **Theorem 34** ([8]). *Consider a reachability bidding game $\mathcal{G}$ played on a general graph. Each vertex v in $\mathcal{G}$ has a* surely-winning threshold *ratio $S\text{THR}(v)$ such that, when Player 1's budget is B and Player 2's budget is 1:*

- *If $B > S\text{THR}(v)$, Player 1 has a deterministic surely-winning strategy.*
- *If $B < S\text{THR}(v)$, Player 2 wins with positive probability.*

Moreover, when the targets are t_1 and t_2, then $S\text{THR}(t_1) = 0$ and $S\text{THR}(t_2) = \infty$,[4] and for every other vertex v, we have $S\text{THR}(v) = S\text{THR}(v^-) + (1 - S\text{THR}(v^-)/S\text{THR}(v^+))$. Finding the thresholds is in PSPACE in general and in linear time in games played on DAGs.

In DAGs, surely-winning threshold ratios are rational numbers. In general graphs, however, Fig. 8 shows that this not the case. The algorithm for computing surely-winning threshold ratios in DAGs is a simple backwards induction. We implemented it and used it to solve the game tic-tac-toe. The solution of tic-tac-toe under first-price bidding is discussed in the blog post `https://bit.ly/2KUong4`: with first-price Richman bidding, Player 1 can guarantee winning when his ratio exceeds $133/123 \approx 1.0183$ [18] and with first-price poorman bidding, when it exceeds roughly 1.0184. Under all-pay poorman bidding, the surely-winning threshold ratio is $51/31 \approx 1.65$ (see Fig. 7). A reason for the slightly higher threshold in all-pay is that contrary to first-price bidding, there is a gap between the surely-winning thresholds of the two players; namely, in the range $(31/51, 51/31)$ neither player surely wins the game. For more unexpected phenomena that the solution points to, see [8].

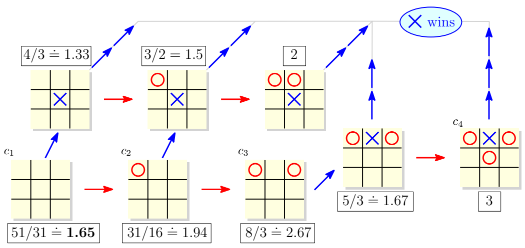

Figure 7 Surely-winning thresholds in all-pay poorman tic-tac-toe. For example, the surely-winning threshold budget in c_4 is 3 since in order to win the game, Player 1 must win three biddings in a row.

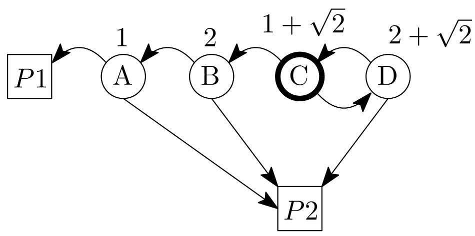

Figure 8 A reachability all-pay poorman game with target Pi for Player i in which surely-winning threshold ratios are irrational.

▶ **Open problem 5.** *We provide a complete picture for the game $G(2, 1)$. Solutions to more complicated games are left open. In particular:*

- *Show that the value exists in general.*
- *Can the value as a function of the budget ratio be a continuous function?*
- *Find closed-form solutions to other race games or other games on DAGs.*

[4] Recall that Player 2's budget is normalized to 1 hence the different values in the targets compared with first-price bidding.

Infinite-duration all-pay bidding games

Recall that even though reachability first-price poorman games do not exhibit an equivalence with random-turn games (Remark 12), an equivalence "pops up" in mean-payoff first-price poorman games (Thm. 20). Inspired by this result, we study the following question:

Are mean-payoff all-pay bidding games equivalent to random-turn games?

This question is studied in [9] and answered positively. Moreover, we find the results particularly surprising. We describe the key properties of mean-payoff games played on strongly connected graphs. Under all-pay Richman bidding, a simple argument shows that deterministic strategies cannot guarantee anything; e.g., in $\mathcal{G}_{\bowtie}$ (Fig. 3), Max cannot deterministically guarantee any positive payoff. On the other hand, with mixed strategies, all-pay and first-price Richman coincide: in both, the initial ratio does not matter, and the optimal *almost-sure* payoff a player can guarantee under all-pay Richman coincides with what he can guarantee deterministically under first-price Richman. In $\mathcal{G}_{\bowtie}$ this is 0.5.

Given these results, it seems safe to guess that poorman all-pay bidding would exhibit similar properties: deterministic strategies are useless and with mixed strategies, first-price and all-pay coincide. Both guesses, however, are wrong. Consider $\mathcal{G}_{\bowtie}$ and suppose Max's ratio is 0.75. As a baseline, recall that under first-price bidding, the mean-payoff value is 0.75. The first surprise is that deterministic strategies are useful under all-pay poorman bidding when the ratio is greater than 0.5; Max can deterministically guarantee a payoff of 2/3 with a ratio of 0.75. The real surprise is with mixed strategies: given a choice between first-price and all-pay bidding, Max strictly prefers all-pay bidding! In $\mathcal{G}_{\bowtie}$, with ratio 0.75, Max can almost-surely guarantee a payoff of 5/6.

5 Discrete Bidding Games

In discrete bidding, the budgets are given in coins, and the minimal positive bid a player can make is one coin. Discrete bidding is more natural with Richman bidding since with poorman bidding, the budgets run out eventually. Motivated by recreational play, reachability discrete bidding games were introduced and studied in [18] (see also [11, 22]). They showed existence of threshold ratios, i.e., an adaption of Thm. 7 to discrete bidding.

Unlike continuous bidding, in discrete bidding, ties in biddings need to be addressed explicitly. We thus assume a bidding game comes equipped with a tie-breaking mechanism. A *configuration* of a bidding game consists of the current vertex the token is placed on, the budgets of the two players, and the state of the tie-breaking mechanism. A game is *determined* if from each configuration exactly one player has a winning strategy. It is not hard to show that bidding games are a sub-class of concurrent games, which are not in general determined. For example, the game "matching pennies" is a non-determined concurrent game: both players concurrently choose a side of a coin, either 1 ("heads") or 0 ("tails"), and Player 1 wins iff the parity of the sum of the players' choices is 0. Matching pennies is not determined since if Player 1 reveals his choice first, Player 2 will choose opposite and win the game, and dually for Player 2.

In [1], we studied the following question:

Which tie-breaking mechanisms admit determinacy of discrete bidding games?

We consider the following classes of tie-breaking mechanisms.

▶ **Definition 35** (Tie-breaking mechanisms)**.**

- ***Transducer-based:*** *Several natural tie-breaking mechanisms can be expressed using a* transducer*. Examples include "Player* 1 *wins all ties", "the players alternate turns in winning ties", "tie breaking depends on the vertex on which the token is placed", "the player with less budget wins ties", and more. Formally, a transducer is* $T = \langle \Sigma, Q, q_0, \Delta, \Gamma\rangle$*, where* Σ *is a set of letters,* Q *is a set of states,* $q_0 \in Q$ *is an initial state,* $\Delta : Q \times \Sigma \to Q$ *is a deterministic transition function, and* $\Gamma : Q \to \{1, 2\}$ *is a labeling of the states. Intuitively,* T *is run in parallel to the bidding game and its state is updated according to the outcomes of the biddings. Whenever a tie occurs and* T *is in state* $s \in Q$*, the winner of the bidding is* $\Gamma(s)$*. A critical point is the information according to which tie-breaking is determined, and this is represented in the alphabet of the transducer.*
- ***Random-based:*** *A tie is resolved by choosing the winner uniformly at random.*
- ***Advantage-based:*** *Exactly one player holds the* advantage*. Suppose Player* i *holds the advantage and a tie occurs. Then Player* i *chooses who wins the bidding. If he calls the other player the winner, Player* i *keeps the advantage, and if he calls himself the winner, the advantage switches to the other player. This is the main tie-breaking mechanism studied in [18].*

The following simple example shows that bidding games are not in general determined with transducer-based tie-breaking. In [1], a more complicated example shows non-determinacy for a particularly simple tie-breaking mechanism: the players alternate turns in winning ties, that is, Player 1 wins the first tie, Player 2 the second tie, Player 1 the third, etc.

▶ **Example 36.** In a *Büchi game*, Player 1 wins a play iff it visits an accepting state infinitely often. Consider the Büchi bidding game that is depicted on the left of Fig. 9, and assume the players' budgets are positive and equal. We claim that the game is not determined when tie-breaking is resolved using the transducer on the right of the figure. That is, if a tie occurs in the first bidding, Player 2 wins all ties for the rest of the game, and otherwise Player 1 wins all ties. First note that, for $i \in \{1, 2\}$, no matter what the budgets are, if Player i wins all ties, he wins the game. A winning strategy for Player 1 always bids 0. Intuitively, Player 2 must invest a unit of budget for winning a bidding and leaving v_1. Thus, if Player 2's initial budget is B, he can leave v_1 only B times, and the game eventually stays in v_1. So, the winner is determined according to the outcome of the first bidding, and the players essentially play a matching-pennies game in that round. ⌟

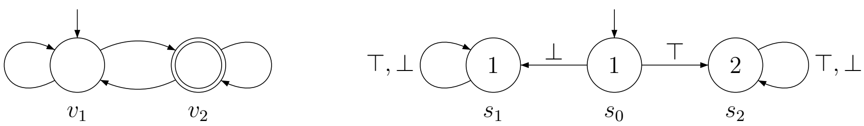

Figure 9 On the left, a Büchi game that is not determined when tie-breaking is determined according to the transducer on the right, where the letters $\top$ and $\bot$ respectively represent "tie" and "no tie".

The statement of the theorem below uses *Müller* objectives, which generalize parity objectives.

▶ **Theorem 37** ([1])**.** *Discrete bidding games are a determined sub-class of concurrent games in the following cases:*

- *Müller bidding games that use a transducer that is not aware of ties are determined.*
- *Reachability bidding games with random tie-breaking are determined.*
- *Müller bidding games with advantage-based tie-breaking are determined.*

Determinacy implies an immediate upper bound on solving discrete bidding games in which the budgets are given in unary.

▶ **Theorem 38** ([1]). *Let $\mathcal{G}$ be a discrete bidding game with objective γ and with either transducer or advantage-based tie-breaking. The complexity of deciding whether Player 1 wins from a given vertex with budgets N_1 and N_2 for the two players, given in unary, matches the complexity of solving a turn-based game with objective γ.*

▶ **Open problem 6.** *Our understanding of infinite-duration discrete bidding games is far from complete. We list specific directions for future work below:*

- *Adapt the theory of infinite-duration continuous bidding games to discrete bidding:*
  - *Recall that the key property of parity continuous-bidding games is that exactly one player wins a strongly-connected game (Lem. 13). This is not true for parity discrete-bidding games. Find a classification for strongly-connected parity discrete-bidding games that can serve a basis for a solution for games played on general graphs.*
  - *Mean-payoff discrete-bidding games were never studied. Can the equivalence with random-turn games in continuous-bidding games be adapted to discrete-bidding?*
- *Discrete bidding is especially suited for practical applications; find concrete applications for this model, e.g., in verifying blockchain technology (see Sec. 4).*
- *Preliminary results on discrete all-pay Richman bidding games were shown in [26] and it is interesting to continue this line of work.*
- *Study the computation complexity of discrete bidding games with budgets given in binary.*

6 Conclusions

We studied three orthogonal dimensions of bidding mechanisms: *who pays* (first-price vs. all-pay), *who is paid* (Richman, poorman, and taxman), and *what types of bids are allowed* (continuous vs. discrete). Our results depict a relatively complete picture, especially for continuous bidding. We believe that there is still ample room for future work on bidding games. In addition to the open problems mentioned throughout the survey, we discuss several orthogonal dimensions to the three mentioned above:

The underlying structure. Bidding games in this survey are always played on graphs. In [6], we established initial results on bidding games that have stochastic behavior: the game is played on an MDP and in each turn, we hold an auction to determine which action is chosen, and the next position of the token is determined according to a probability distribution that depends on the current vertex and the chosen action. Other possible structures to consider include real-time systems or systems with counters.

Objectives. We focused on three objectives: reachability, parity, and mean-payoff. It is interesting to study bidding games with other objectives. For example, in *discounted-sum* bidding games, the players are encouraged to bid higher in the first rounds due to the discounted payoffs. Another example is *quantitative reachability*, which is played on a weighted directed graph, the game ends once a target is reached, and the payoff is the (un-averaged) sum of the traversed weights.

One motivation to study bidding games is that they model ongoing auctions. In auctions, and different from mean-payoff objectives, the budget has a value to the players. We thus find it interesting to incorporate the budget in the payoff. For example, suppose that a play in a quantitative reachability game, reaches the target after traversing weights with a sum of

c, and leaving Player 1 with a budget of B. Then, we can define Player 1's payoff as $c + B$. That is, he gains a utility of c for the items he purchased and a utility of B from his left-over money. A similar definition can be applied to infinite plays like in mean-payoff games.

Non-zero-sum games. To the best of our knowledge, non-zero-sum bidding games were only studied in [25]. They consider two-player games on DAGs in which each leaf is labeled by two (not necessarily contradicting) payoffs that the players gain when the game ends in the leaf. They show existence of *subgame perfect equilibrium* in bidding games with out-degree 2. We believe that their results encourage further study of non-zero-sum bidding games.

Multi-player games. To the best of our knowledge, a combination of bidding games with multiple players has never been considered. While it is tempting to study non-zero-sum multi-player games, we note that it is interesting to study a zero-sum version in the spirit of [2]. For example, consider a reachability game with three players, where Player 1 has a target and Players 2 and 3 win iff Player 1 does not reach his target. We hold a first-price bidding in each round to determine who moves the token (say poorman bidding). The coalition can coordinate their actions, but they are restricted compared with a "traditional" player in a two-player bidding game: For example, if the budgets are $\langle 1.5, 1, 1\rangle$, even though the sum of budgets in the coalition is 2, the highest a member of the coalition can bid is 1, thus Player 1 can guarantee winning the bidding.

We mention a surprising result regarding this model. We say that $\overline{B} = \langle B_1, B_2, B_3\rangle \in \mathbb{R}^3$ is a *winning point* if Player 1 wins when Player i's budget is B_i, for $i \in \{1, 2, 3\}$. Ventsislav Chonev showed that the set $\{\overline{B} : \overline{B} \text{ is a winning point}\}$ is not convex in the game in which the coalition needs to win two biddings and Player 1 needs to win once.

Changes to the bidding mechanism. We list two examples of changes to the bidding mechanism, and there are plenty of other interesting changes one can study. It is appealing, especially in discrete poorman bidding, to add re-charging to the budgets. One way recharging can be defined is by labeling each vertex with a pair $\langle c_1, c_2\rangle \in \mathbb{N}^2$ and when the token lands on a vertex, we add c_i to Player i's budget, for $i \in \{1, 2\}$. So, Player i's available budget at the end of a finite play is his initial budget, minus the sum of bids he pays, plus the sum of recharges to his budget.

A second change to the bidding mechanism relaxes the requirement that the bidding mechanism applies in the same manner to all vertices in the graph. A simple definition would be to partition the vertices into vertices in which we hold a Richman bidding and those in which we hold a poorman bidding. For example, consider the Bowtie game (Fig. 3) and set one of the vertices to be Richman and the other to poorman. It is not clear whether Max prefers his vertex being the poorman vertex, the Richman vertex, or whether he has no preference. More generally, each vertex can be associated with its own taxman parameter.

References

1 M. Aghajohari, G. Avni, and T. A. Henzinger. Determinacy in discrete-bidding infinite-duration games. In *Proc. 30th CONCUR*, volume 140 of *LIPIcs*, pages 20:1–20:17, 2019.

2 R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.

3 K.R. Apt and E. Grädel. *Lectures in Game Theory for Computer Scientists.* Cambridge University Press, 2011.

4 G. Avni, T. A. Henzinger, and V. Chonev. Infinite-duration bidding games. *J. ACM*, 66(4):31:1–31:29, 2019.

5 G. Avni, T. A. Henzinger, and R. Ibsen-Jensen. Infinite-duration poorman-bidding games. In *Proc. 14th WINE*, volume 11316 of *LNCS*, pages 21–36. Springer, 2018.

6 G. Avni, T. A. Henzinger, R. Ibsen-Jensen, and P. Novotný. Bidding games on markov decision processes. In *Proc. 13th RP*, pages 1–12, 2019.

7 G. Avni, T. A. Henzinger, and Đ. Žikelić. Bidding mechanisms in graph games. In *In Proc. 44th MFCS*, volume 138 of *LIPIcs*, pages 11:1–11:13, 2019.

8 G. Avni, R. Ibsen-Jensen, and J. Tkadlec. All-pay bidding games on graphs. *Proc. 34th AAAI*, 2020.

9 G. Avni, I. Jecker, and Đ. Žikelić. Infinite-duration all-pay bidding games. *CoRR*, abs/2005.06636, 2020. `arXiv:2005.06636`.

10 M. R. Baye and H. C. Hoppe. The strategic equivalence of rent-seeking, innovation, and patent-race games. *Games and Economic Behavior*, 44(2):217–226, 2003.

11 J. Bhatt and S. Payne. Bidding chess. *Math. Intelligencer*, 31:37–39, 2009.

12 E. Borel. La théorie du jeu les équations intégrales á noyau symétrique. *Comptes Rendus de l'Académie*, 173(1304–1308):58, 1921.

13 T. Brázdil, V. Brozek, K. Etessami, and A. Kucera. Approximating the termination value of one-counter mdps and stochastic games. In *Proc. 38th ICALP*, pages 332–343, 2011.

14 T. Brázdil, V. Brozek, K. Etessami, A. Kucera, and D. Wojtczak. One-counter markov decision processes. In *Proc. 21st SODA*, pages 863–874, 2010.

15 J. F. Canny. Some algebraic and geometric computations in PSPACE. In *Proc. 20th STOC*, pages 460–467, 1988.

16 K. Chatterjee, A. K. Goharshady, and Y. Velner. Quantitative analysis of smart contracts. In *Proc. 27th ESOP*, pages 739–767, 2018.

17 A. Condon. The complexity of stochastic games. *Inf. Comput.*, 96(2):203–224, 1992.

18 M. Develin and S. Payne. Discrete bidding games. *The Electronic Journal of Combinatorics*, 17(1):R85, 2010.

19 A. E. Emerson, C. S. Jutla, and P. A. Sistla. On model-checking for fragments of μ-calculus. In *Proc. 5th CAV*, pages 385–396, 1993.

20 A. R. Howard. *Dynamic Programming and Markov Processes*. MIT Press, 1960.

21 K. A. Konrad and D. Kovenock. Multi-battle contests. *Games and Economic Behavior*, 66(1):256–274, 2009.

22 U. Larsson and J. Wästlund. Endgames in bidding chess. *Games of No Chance 5*, 70, 2018.

23 A. J. Lazarus, D. E. Loeb, J. G. Propp, W. R. Stromquist, and D. H. Ullman. Combinatorial games under auction play. *Games and Economic Behavior*, 27(2):229–264, 1999.

24 A. J. Lazarus, D. E. Loeb, J. G. Propp, and D. Ullman. Richman games. *Games of No Chance*, 29:439–449, 1996.

25 R. Meir, G. Kalai, and M. Tennenholtz. Bidding games and efficient allocations. *Games and Economic Behavior*, 2018. `doi:10.1016/j.geb.2018.08.005`.

26 M. Menz, J. Wang, and J. Xie. Discrete all-pay bidding games. *CoRR*, abs/1504.02799, 2015.

27 Y. Peres, O. Schramm, S. Sheffield, and D. B. Wilson. Tug-of-war and the infinity laplacian. *J. Amer. Math. Soc.*, 22:167–210, 2009.

28 Y. Peres and Z. Sunic. Biased infinity laplacian boundary problem on finite graphs. *CoRR*, abs/1912.13394, 2019. `arXiv:1912.13394`.

29 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.

30 M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 2005.

31 M.O. Rabin. Decidability of second order theories and automata on infinite trees. *Transaction of the AMS*, 141:1–35, 1969.

32 G. Tullock. *Toward a Theory of the Rent Seeking Society*, chapter Efficient rent seeking, pages 97–112. College Station: Texas A&M Press, 1980.