Bidding Graph Games with Partially-Observable Budgets

Guy Avni[1], Ismäel Jecker[2], Đorđe Žikelić[3]

[1]University of Haifa
[2]University of Warsaw
[3]Institute of Science and Technology Austria (ISTA)

Abstract

Two-player zero-sum *graph games* are a central model, which proceeds as follows. A token is placed on a vertex of a graph, and the two players move it to produce an infinite *play*, which determines the winner or payoff of the game. Traditionally, the players alternate turns in moving the token. In *bidding games*, however, the players have budgets and in each turn, an auction (bidding) determines which player moves the token. So far, bidding games have only been studied as full-information games. In this work we initiate the study of partial-information bidding games: we study bidding games in which a player's initial budget is drawn from a known probability distribution. We show that while for some bidding mechanisms and objectives, it is straightforward to adapt the results from the full-information setting to the partial-information setting, for others, the analysis is significantly more challenging, requires new techniques, and gives rise to interesting results. Specifically, we study games with *mean-payoff* objectives in combination with *poorman* bidding. We construct optimal strategies for a partially-informed player who plays against a fully-informed adversary. We show that, somewhat surprisingly, the *value* under pure strategies does not necessarily exist in such games.

Introduction

We consider two-player zero-sum *graph games*; a fundamental model with applications, e.g., in multi-agent systems (Alur, Henzinger, and Kupferman 2002). A graph game is played on a finite directed graph as follows. A token is placed on a vertex and the players move it throughout the graph to produce an infinite path, which determines the payoff of the game. Traditional graph games are *turn-based*: the players alternate turns in moving the token.

Bidding games (Lazarus et al. 1996, 1999) are graph games in which an "auction" (bidding) determines which player moves the token in each turn. The concrete bidding mechanisms that we consider proceed as follows. In each turn, both players simultaneously submit bids, where a bid is legal if it does not exceed the available budget. The higher bidder "wins" the bidding and moves the token. The mechanisms differ in their payment schemes, which are classified according to two orthogonal properties. *Who pays:* in

first-price bidding only the higher bidder pays the bid and in *all-pay* bidding both players pay their bids. *Who is the recipient:* in *Richman* bidding (named after David Richman) payments are made to the other player and in *poorman* bidding payments are made to the "bank", i.e., the bid is lost. As a rule of thumb, bidding games under all-pay and poorman bidding are respectively technically more challenging than first-price and Richman bidding. More on this later. In terms of applications, however, we argue below that poorman bidding is often the more appropriate bidding mechanism.

Applications. A central application of graph games is *reactive synthesis* (Pnueli and Rosner 1989): given a specification, the goal is to construct a *controller* that ensures correct behavior in an adversarial environment. Synthesis is solved by constructing a turn-based *parity* game in which Player 1 is associated with the controller and Player 2 with the environment, and searching for a winning Player 1 strategy.

Bidding games extend the modeling capabilities of graph games. For example, they model ongoing and stateful auctions in which budgets do not contribute to the players' utilities. Advertising campaigns are one such setting: the goal is to maximize visibility using a pre-allocated advertising budget. By modeling this setting as a bidding game and solving for Player 1, we obtain a bidding strategy with guarantees against any opponent[1]. Maximizing visibility can be expressed as a mean-payoff objective (defined below).

All-pay poorman bidding is particularly appealing since it constitutes a dynamic version of the well-known Colonel Blotto games (Borel 1921). Rather than thinking of the budgets as money, we think of them as resources at the disposal of the players, like time or energy. Then, deciding how much to bid represents the effort that a player invests in a competition, e.g., investing time to prepare for a job interview, where the player that invests more wins the competition.

Prior work – full-information bidding games. The central quantity in bidding games is the *initial ratio* between the players' budgets. Formally, for $i \in \{1, 2\}$, let B_i be Player i's initial budget. Then, Player 1's initial ratio is $B_1/(B_1 + B_2)$. A *random-turn game* (Peres et al. 2009) with parameter $p \in [0, 1]$ is similar to a bidding game only that in-

---

[1]A worst-case modelling assumes that the other bidders cooperate against Player 1.

stead of bidding, in each turn, we toss a coin with probability p that determines which player moves the token. Formally, a random-turn game is a special case of a *stochastic game* (Condon 1992).

Qualitative objectives. In *reachability* games, each player is associated with a target vertex, the game ends once a target is reached, and the winner is the player whose target is reached. Reachability bidding games were studied in (Lazarus et al. 1996, 1999). It was shown that, for first-price reachability games, a *threshold ratio* exists, which, informally, is a necessary and sufficient initial ratio for winning the game. Moreover, it was shown that first-price Richman-bidding games are equivalent to uniform random-turn games (and only Richman bidding); namely, the threshold ratio in a bidding game corresponds to the value of a uniform random-turn game. All-pay reachability games are technically more challenging. Optimal strategies might be mixed and may require sampling from infinite-support probability distributions even in extremely simple games (Avni, Ibsen-Jensen, and Tkadlec 2020).

Mean-payoff games. Mean-payoff games are infinite-duration quantitative games. Technically, each vertex of the graph is assigned a weight, and the *payoff* of an infinite play is the long-run average sum of weights along the path. The payoff is Player 1's reward and Player 2's cost, thus we refer to them respectively as Max and Min. For example, consider the "bowtie" game $\mathcal{G}_{\bowtie}$, depicted in Fig. 1. The payoff in $\mathcal{G}_{\bowtie}$ corresponds to the ratio of bidding that Max wins. Informally $\mathcal{G}_{\bowtie}$ models the setting in which in each day a publisher sells an ad slot, and Max's objective is to maximize visibility: the number of days that his ad is displayed throughout the year. Unlike reachability games, intricate equivalences between mean-payoff bidding games and random-turn games are known for all the mechanisms described above (Avni, Henzinger, and Chonev 2019; Avni, Henzinger, and Ibsen-Jensen 2018; Avni, Henzinger, and Žikelić 2021; Avni, Jecker, and Žikelić 2021).

Example 1. We illustrate the equivalences between full-information bidding games and random-turn games. Consider the "bowtie" game $\mathcal{G}_{\bowtie}$ (see Fig. 1). For $p \in [0,1]$, the random-turn game $\mathsf{RT}(\mathcal{G}_{\bowtie}, p)$ that uses a coin with bias p is depicted in Fig. 2. Its expected payoff is p.

Suppose that the initial ratio is $r \in (0,1)$. Under first-price Richman-bidding, the optimal payoff in $\mathcal{G}_{\bowtie}$ does not depend on the initial ratio: no matter what r is, the optimal payoff that Max can guarantee is arbitrarily close to 0.5, hence the equivalance with $\mathsf{RT}(\mathcal{G}_{\bowtie}, 0.5)$. Under first-price poorman bidding, the optimal payoff *does* depend on the initial ratio: roughly, the optimal payoff that Max can guarantee is r, hence the equivalence with $\mathsf{RT}(\mathcal{G}_{\bowtie}, r)$. For all-pay bidding, pure strategies are only "useful" in all-pay poorman bidding and only when $r > 0.5$, where Max can guarantee an optimal payoff of $\frac{2r-1}{r}$. The results extend to general strongly-connected games (see Thm. 3). △

Our contributions – partial-information bidding games. In most auction domains, bidders are not precisely informed of their opponent's budget. Bidding games, however, have only been studied as full-information games. We initiate the
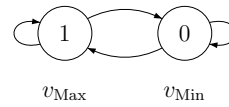

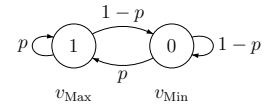Figure 1: The mean-payoff game $\mathcal{G}_{\bowtie}$ with the weights in the vertices.

Figure 2: The simplified random-turn game $\mathsf{RT}(\mathcal{G}_{\bowtie}, p)$, for $p \in [0,1]$.

study of bidding games in which the players are partially informed of the opponent's budget. Specifically, we study bidding games in which the two players' budgets are drawn from a known probability distribution, and the players' goal is to maximize their expected utility. We first show that the results on qualitative objectives as well as first-price Richman bidding transfer to the partial-information setting.

We turn to study mean-payoff poorman-bidding games, which are significantly more challenging. We focus on *one-sided* partial-information games in which only Player 2's budget is drawn from a probability distribution. Thus, Player 1 is *partially informed* and Player 2 is *fully informed* of the opponent's budget. We argue that one-sided partial-information games are practically well-motivated. Indeed, one-sided partial information is a worst-case modelling: the utility that an optimal strategy for Player 1 guarantees in the game, is a lower bound on the utility that it will guarantee when deployed against the concrete environment. We illustrate our results in the following example.

Example 2. Consider the bowtie game $\mathcal{G}_{\bowtie}$ (Fig. 1), where Max (the partially-informed player) starts with a budget of B and Min (the fully-informed player) starts with a budget that is drawn uniformly at random from $\mathrm{supp}(\gamma) = \{C_1, C_2\}$. We describe an optimal strategy for Max under first-price poorman bidding. Max carefully chooses an $x \in [B \cdot \frac{C_1}{C_2}, B]$ and divides his budget into two "wallets"; the first with budget x and the second with budget $B - x$. He initially uses his first wallet to play an optimal full-information strategy assuming the initial budgets are x and C_1, which guarantees a payoff of at least $p_1 = \frac{x}{C_1 + x}$. If Player 2 spends more than C_1, i.e., her initial budget was in fact C_2, then Player 1 proceeds to use his second wallet against Player 2's remaining budget, which guarantees a payoff of at least $p_2 = \frac{B - x}{B - x + C_2 - C_1}$. Thus, the expected payoff is at least $0.5 \cdot (p_1 + p_2)$, and Max simply chooses an x that maximizes this expression. Note that the constraint that $x \geq B \cdot \frac{C_1}{C_2}$ implies that $p_1 \geq p_2$, thus Min has an incentive to play so that Max proceeds to use his second wallet. We show that this strategy is optimal, and extend the technique to obtain optimal strategies in general strongly-connected games for first-price and all-pay poorman bidding.

Finally, we show that the optimal payoff that Min can guarantee in $\mathcal{G}_{\bowtie}$, is obtained by a surprisingly simple strategy. We show that the following Min strategy is optimal: when her initial budget is C_i, for $i \in \{1, 2\}$, Min follows an optimal full-information strategy for ratio $B/(B + C_i)$. That is, she "reveals" her true budget in the first round and cannot gain utility by hiding this information. The technical challenge is to show that this strategy is optimal. △

Our results show that contrary to turn-based, stochastic games, and full-information bidding games, there is a gap between the optimal payoffs that the players can guarantee with pure strategies. Thus, the *value* does not necessary exist in partial-information mean-payoff bidding games under pure strategies.

Related work. The seminar book (Aumann, Maschler, and Stearns 1995) studies the mean-payoff game $\mathcal{G}_{\bowtie}$ under one-sided partial-information with a different semantic to the one we study. Let L or R denote the two vertices of $\mathcal{G}_{\bowtie}$. Min has partial information of the weights of L and R, which, before the game begins, are drawn from a known probability distribution. Max, the fully-informed player, knows the weights. In each turn, Max chooses L or R, followed by Min who either "accepts" or "rejects" Max's choice, thus both players can affect the movement of the token. The value in the game is shown to exist. Interestingly and similar in spirit to our results, there are cases in which Max cannot use his knowledge advantage and his optimal strategy reveals which of the two vertices he prefers. One-sided partial information have also been considered in turn-based graph games, e.g., (Reif 1984; Raskin et al. 2007; Wulf, Doyen, and Raskin 2006).

Discrete bidding games were studied in (Develin and Payne 2010); namely, budgets are given in coins, and the minimal positive bid a player can make is a single coin. Tie-breaking is a significant factor in such games (Aghajohari, Avni, and Henzinger 2021). Non-zero-sum bidding games were studied in (Meir, Kalai, and Tennenholtz 2018). See also the survey (Avni and Henzinger 2020).

Preliminaries

Strategies in bidding games. A bidding game is played on a directed graph $\langle V, E \rangle$. A *strategy* in any graph game is a function from histories to actions. In bidding games, a history consists of the sequence of vertices that were visited and bids made by the two players. We stress that the history does not contain the current state of the budgets. Rather, a player can compute his opponent's current budget based on the history of bids, if he knows her initial budget. We formalize the *available budget* following a history. For $i \in \{1, 2\}$, suppose the initial budget of Player i is B_i. For a history h, we define the *investments* of Player i throughout h, denoted $\mathrm{Inv}_i(h)$. In all-pay bidding, $\mathrm{Inv}_i(h)$ is the sum bids made by Player i throughout h, and in first-price bidding, it is the sum only over the winning bids. We denote by $B_i(h)$ Player i's available budget following h. Under Richman bidding, winning bids are paid to the opponent, thus $B_i(h) = B_i - \mathrm{Inv}_i(h) + \mathrm{Inv}_{3-i}(h)$. Under poorman bidding, winning bids are paid to the bank, thus $B_i(h) = B_i - \mathrm{Inv}_i(h)$.

Given a history, a strategy prescribes an action, which in a bidding game, is a pair $\langle b, u \rangle \in \mathbb{R} \times V$, where b is a bid and u is the vertex to move to upon winning. We restrict the actions of the players following a history h so that (1) the bid does not exceed the available budget, thus following a history h, a legal bid for Player i is a bid in $[0, B_i(h)]$, and (2) a player must choose a neighbor of the vertex that

the token is placed on. We restrict attention to strategies that choose legal actions for all histories. Note that we consider only *pure* strategies and disallow *mixed* strategies (strategies that allow a random choice of action).

Definition 1. *For $i \in \{1, 2\}$, we denote by $\mathcal{S}_i(B_i)$ the set of legal strategies for Player i with an initial budget of B_i. Note that with a higher initial budget, there are more strategies to choose from, i.e., for $B_i' > B_i$, we have $\mathcal{S}_i(B_i) \subseteq \mathcal{S}_i(B_i')$.*

The central quantity in bidding games is the initial ratio, defined as follows.

Definition 2. Budget ratio. *When Player i's budget is B_i, for $i \in \{1, 2\}$, we say that Player i's ratio is $\frac{B_i}{B_1 + B_2}$.*

Plays. Consider initial budgets B_1 and B_2 for the two players, two strategies $f \in \mathcal{S}_1(B_1)$ and $g \in \mathcal{S}_2(B_2)$, and an initial vertex v. The triple f, g, and v gives rise to a unique *play*, denoted $\mathsf{play}(v, f, g)$. The construction of $\mathsf{play}(v, f, g)$ is inductive and is intuitively obtained by allowing the players to play according to f and g. Initially, we place the token on v, thus the first history of the game is $h = v$. Suppose a history h has been played. Then, the next action that the players choose is respectively $\langle u_1, b_1 \rangle = f(h)$ and $\langle u_2, b_2 \rangle = g(h)$. If $b_1 > b_2$, then Player 1 wins the bidding and the token moves to u_1, and otherwise Player 2 wins the bidding and the token moves to u_2. Note that we resolve ties arbitrarily in favor of Player 2. The play continues indefinitely. Since the players always choose neighboring vertices, each play corresponds to an infinite path in $\langle V, E \rangle$. For $n \in \mathbb{N}$, we use $\mathsf{play}_n(v, f, g)$ to denote its finite prefix of length n. We sometimes omit the initial vertex from the play when it is clear from the context.

Objectives. We consider zero-sum games. An *objective* assigns a *payoff* to a play, which can be thought of as Player 1's reward and Player 2's penalty. We thus sometimes refer to Player 1 as Max and Player 2 as Min. We denote by $\mathsf{payoff}(f, g, v)$ the payoff of the play $\mathsf{play}(f, g, v)$.

Qualitative objectives. The payoff in games with qualitative objectives is in $\{-1, 1\}$. We say that Player 1 *wins* the play when the payoff is 1. We consider two qualitative objectives. (1) *Reachability.* There is a distinguished target vertex t and a play is winning for Player 1 iff it visits t. (2) *Parity.* Each vertex is labeled by an index in $\{1, \ldots, d\}$ and a play is winning for Player 1 iff the highest index that is Parity objectives are important in practice, e.g., reactive synthesis (Pnueli and Rosner 1989) is reduced to the problem of solving a (turn-based) parity games.

Mean-payoff games. The quantitative objective that we consider is *mean-payoff*. Every vertex v in a mean-payoff game has a weight $w(v)$ and the payoff of an infinite play is the long-run average weight that it traverses. Formally, the payoff of an infinite path $v_1, v_2, \ldots$ is $\liminf_{n \to \infty} \frac{1}{n} \sum_{1 \le i < n} w(v_i)$. Note that the definition favors Min since it uses $\liminf$.

Values in full-information bidding games. We are interested in finding the optimal payoff that a player can *guarantee* with respect to an initial budget ratio. Let $c \in \mathbb{R}$ and

initial budgets B_1 and B_2. We say that Player 1 can *guarantee* a payoff of c, if he can reveal that he will be playing according to a strategy $f \in \mathcal{S}_1(B_1)$, and no matter which strategy $g \in \mathcal{S}_2(B_2)$ Player 2 responds with, we have payoff$(f, g) \geq c$. *Player 1's value* is the maximal c that he can guarantee, and Player 2's value is defined dually. Note that there might be a gap between the two players' values. When Player 1's value coincides with Player 2's value, we say that the *value* exists in the game.

Partial Information Bidding Games

A partial-information bidding game is $\mathcal{G} = \langle V, E, \alpha, \gamma_1, \gamma_2 \rangle$, where $\langle V, E \rangle$ is a directed graph, α is an objective as we elaborate later, and the *budget distribution* γ_i is a probability distribution from which Player i's initial budget is drawn, for $i \in \{1, 2\}$. The *support* of a probability distribution $\gamma : \mathbb{Q} \rightarrow [0, 1]$ is supp$(\gamma) = \{x \in \mathbb{Q} : \gamma(x) > 0\}$. We restrict attention to finite-support probability distributions. For $i \in \{1, 2\}$, the probability that Player i's initial budget is $B_i \in$ supp(γ_i) is $\gamma_i(B_i)$.

Definition 3. One-sided partial information. *We say that a game has one-sided partial information when $|supp(\gamma_1)| = 1$ and $|supp(\gamma_2)| > 1$. We then call Player 1 the* partially-informed player *and Player 2 the* fully-informed player.

We turn to define values in partial-information games. The intuition is similar to the full-information case only that each player selects a collection of strategies, one for each possible initial budget, and we take the expectation over the payoffs that each pair of strategies achieves. The δ in the following definition allows us to avoid corner cases due to ties in biddings and the ϵ is crucial to obtain the results on full-information mean-payoff bidding games.

Definition 4. (Values in partial-information bidding games). *Consider a partial-information bidding game $\mathcal{G} = \langle V, E, \alpha, \beta, \gamma \rangle$. Suppose $supp(\beta) = \{B_1, \ldots, B_n\}$ and $supp(\gamma) = \{C_1, \ldots, C_n\}$. We define Player 1's value, denoted $\mathsf{val}^{\downarrow}(\mathcal{G}, \beta, \gamma)$, and Player 2's value, denoted $\mathsf{val}^{\uparrow}(\mathcal{G}, \beta, \gamma)$, is defined symmetrically. We define that $\mathsf{val}^{\downarrow}(\mathcal{G}, \beta, \gamma) = c \in \mathbb{R}$ if for every $\delta, \epsilon > 0$,*

- *There is a collection $\left(f_B \in \mathcal{S}_1(B + \delta) \right)_{B \in supp(\beta)}$ of Player 1 strategies, such that for every collection $\left(g_C \in \mathcal{S}_2(C) \right)_{C \in supp(\gamma)}$ of Player 2 strategies, we have $\sum_{B,C} \beta(B) \cdot \gamma(C) \cdot \mathsf{payoff}(f_B, g_C) \geq c - \epsilon$.*

- *For every collection $\left(f_B \in \mathcal{S}_1(B) \right)_{B \in supp(\beta)}$ of Player 1 strategies, there is a collection $\left(g_C \in \mathcal{S}_2(C + \delta) \right)_{C \in supp(\gamma)}$ of Player 2 strategies such that $\sum_{B,C} \beta(B) \cdot \gamma(C) \cdot \mathsf{payoff}(f_B, g_C) \leq c + \epsilon$.*

Note that $\mathsf{val}^{\downarrow}(\mathcal{G}, \beta, \gamma) \leq \mathsf{val}^{\uparrow}(\mathcal{G}, \beta, \gamma)$ and when there is equality, we say that the value exists, and denote it by $\mathsf{val}(\mathcal{G}, \beta, \gamma)$.

The value in mean-payoff games is often called the *mean-payoff value*. In mean-payoff games we use $\mathsf{MP}^{\downarrow}, \mathsf{MP}^{\uparrow}$, and MP instead of $\mathsf{val}^{\downarrow}, \mathsf{val}^{\uparrow}$, and val, respectively. When $\mathcal{G}$ is

full-information and the budget ratio is r, we use $\mathsf{MP}(\mathcal{G}, r)$ instead of writing the two budgets.

Partial-Information Qualitative First-Price Bidding Games

In this section, we focus on first-price bidding and show that the value exists in partial-information bidding games with qualitative objectives. The proof adapts results from the full-information setting, which we survey first.

Definition 5. (Threshold ratios in full-information games). *Consider a full-information first-price bidding game with a qualitative objective. Suppose that the sum of initial budgets is 1 and that the game starts at v. The threshold ratio in v, denoted $Th(v)$, is a value t such that for every $\epsilon > 0$:*

- *Player 1 wins when his ratio is greater than $Th(v)$; namely, when the initial budgets are $t + \epsilon$ and $1 - t - \epsilon$.*
- *Player 2 wins when Player 1's ratio is less than $Th(v)$; namely, when the initial budgets are $t - \epsilon$ and $1 - t + \epsilon$.*

Existence of threshold ratios for full-information reachability games was shown in (Lazarus et al. 1996, 1999) and later extended to full-information parity games in (Avni, Henzinger, and Chonev 2019; Avni, Henzinger, and Ibsen-Jensen 2018).

Theorem 1. *(Lazarus et al. 1996, 1999; Avni, Henzinger, and Chonev 2019; Avni, Henzinger, and Ibsen-Jensen 2018) Threshold ratios exist in every vertex of a parity game.*

The following theorem, whose proof can be found in the full version, extends these results to the partial-information setting.

Theorem 2. *Consider a partial-information parity first-price bidding game $\mathcal{G} = \langle V, E, \alpha, \beta, \gamma \rangle$ and a vertex $v \in V$. Let $W = \{\langle B, C \rangle : B \in supp(\beta), C \in supp(\gamma), \text{ and } Th(v) < \frac{B}{B+C}\}$. Then, the value of $\mathcal{G}$ in v is $\sum_{\langle B, C \rangle \in W} \beta(B) \cdot \gamma(C)$.*

Partial-Information Mean-Payoff Bidding Games

In this section we study mean-payoff bidding games. Throughout this section we focus on games played on *strongly-connected graphs*. We start by surveying results on full-information games. The most technically-challenging results concern one-sided partial-information poorman-bidding games. We first develop optimal strategies for the partially-informed player, and then show that the value does not necessary exist under pure strategies.

Full-Information Mean-Payoff Bidding Games

We show equivalences between bidding games and a class of stochastic games (Condon 1992) called random-turn games, which are define formally as follows.

Definition 6. (Random-turn games). *Consider a strongly-connected mean-payoff bidding game $\mathcal{G}$. For $p \in [0, 1]$, the random-turn game that corresponds to $\mathcal{G}$ w.r.t. p, denoted $RT(\mathcal{G}, p)$, is a game in which instead of bidding, in*

each turn, we toss a (biased) coin to determine which player moves the token: Player 1 and Player 2 are respectively chosen with probability p and $1-p$. Formally, $\mathsf{RT}(\mathcal{G},p)$ is constructed as follows. Every vertex v in $\mathcal{G}$, is replaced by three vertices v_N, v_1, and v_2. The vertex v_N simulates the coin toss: it has an outgoing edge with probability p to v_1 and an edge with probability $1-p$ to v_2. For $i \in \{1,2\}$, vertex v_i simulates Player i winning the coin toss: it is controlled by Player i and has an outgoing edge to u_N, for every neighbor u of v. The weights of v_N, v_1, and v_2 coincide with the weight of v. The mean-payoff value of $\mathsf{RT}(\mathcal{G},p)$, denoted $\mathsf{MP}(\mathsf{RT}(\mathcal{G},p))$, is the optimal expected payoff that the two players can guarantee, and it is known to exist (Puterman 2005). Since $\mathcal{G}$ is strongly-connected, $\mathsf{MP}(\mathsf{RT}(\mathcal{G},p))$ does not depend on the initial vertex.

For a full-information game $\mathcal{G}$ and a ratio $r \in (0,1)$, recall that $\mathsf{MP}(\mathcal{G},r)$ denotes the optimal payoff that Max can guarantee with initial ratio r. We state the equivalences between the two models.

Theorem 3. *Let $\mathcal{G}$ be a strongly-connected full-information mean-payoff bidding game.*

- **First-price Richman bidding (Avni, Henzinger, and Chonev 2019).** *The optimal payoff that Max can guarantee with a pure strategy does not depend on the initial ratio: for every initial ratio r, we have $\mathsf{MP}(\mathcal{G},r) = \mathsf{MP}(\mathsf{RT}(\mathcal{G},0.5))$.*

- **First-price poorman bidding (Avni, Henzinger, and Ibsen-Jensen 2018).** *The optimal payoff that Max can guarantee with pure strategy and ratio r coincides with the value of a random-turn game with bias r: for every initial ratio r, we have $\mathsf{MP}(\mathcal{G},r) = \mathsf{MP}(\mathsf{RT}(\mathcal{G},r))$.*

- **All-pay poorman bidding (Avni, Jecker, and Žikelić 2021).** *The optimal payoff that Max can guarantee with a pure strategy and ratio $r > 0.5$ coincides with the value of a random-turn game with bias $(2r-1)/r$: for every initial ratio $r > 0.5$, we have $\mathsf{MP}(\mathcal{G},r) = \mathsf{MP}(\mathsf{RT}(\mathcal{G},(2r-1)/r))$.*

Since the optimal payoff under first-price Richman bidding depends only on the structure of the game and not on the initial ratios, the result easily generalizes to partial-information games. Consider two budget distributions β and γ for Min and Max, respectively. Indeed, when Min's initial budget is $B \in \mathrm{supp}(\beta)$, playing optimally against any $C \in \mathrm{supp}(\gamma)$ results in the same payoff, and similarly for Max. We thus conclude the following.

Theorem 4. *Consider a strongly-connected first-price Richman mean-payoff bidding game $\mathcal{G}$. For any two budget distributions β and γ for the two players, we have $\mathsf{MP}^{\downarrow}(\mathcal{G},\beta,\gamma) = \mathsf{MP}^{\uparrow}(\mathcal{G},\beta,\gamma) = \mathsf{MP}(\mathsf{RT}(\mathcal{G},0.5))$.*

Remark 1. (All-pay Richman bidding). It was shown in (Avni, Jecker, and Žikelić 2021) that in all-pay Richman bidding games, pure strategies are "useless": no matter what the initial ratio is, Max cannot guarantee a positive payoff with a pure strategy. The study of mean-payoff all-pay Richman-bidding games is thus trivial in the partial-information setting as well.

The Value of the Partially-Informed Player

We turn to study partial-information mean-payoff bidding games under poorman bidding, where we focus on one-sided partial information. We arbitrarily set Max to be partially-informed and Min to be fully-informed.

First-price bidding. Fix a strongly-connected mean-payoff game $\mathcal{G}$. Suppose that Max's budget is B and Min's budget is chosen from a finite probability distribution γ with $\mathrm{supp}(\gamma) = \{C_1,\ldots,C_n\}$ and $C_i < C_{i+1}$, for $1 \leq i < n$. We generalize the technique that is illustrated in Example 2. Max carefully chooses increasing $x_1,\ldots,x_n$, where $x_n = B$. He maintains two "accounts": a spending account from which he bids and a savings account. Initially, the spending account has a budget of x_1 and the savings account, a budget of $B - x_1$. Max plays "optimistically". He first plays in hope that Min's budget is C_1 with a budget of x_1. If Min does not spend C_1, the payoff is as in full-information games, namely at least $p_1 = \mathsf{MP}(\mathsf{RT}(\mathcal{G},\frac{x_1}{x_1+C_1}))$. Otherwise, Min spends at least C_1 and Max transfers budget from his savings account to his spending account so that the saving account has $B - x_2$ and the spending account has at least $x_2 - x_1$. Note that if Min's initial budget was indeed C_2, at this point she is left with a budget of at most $C_2 - C_1$. If Min does not spend $C_2 - C_1$, by following a full-information optimal strategy, Max can guarantee a payoff of at least $p_2 = \mathsf{MP}(\mathsf{RT}(\mathcal{G},\frac{x_2-x_1}{x_2-x_1+C_2-C_1}))$. The definition of $p_3,\ldots,p_n$ is similar. Max chooses $x_1,\ldots,x_n$ so that $p_1 \geq \ldots \geq p_n$. Thus, when Min's initial budget is C_i, she has an incentive to play so that Max's spending account will reach x_i and the payoff will be at least p_i. We call such a choice of $x_1,\ldots,x_n$ *admissible* and formally define it as follows.

Definition 7. Admissible sequences. *Let $\mathcal{G}$ be a poorman mean-payoff bidding game. Let B be a budget of Max and γ be a finite budget distribution of Min with $\mathrm{supp}(\gamma) = \{C_1,\ldots,C_n\}$. A sequence $(x_i)_{1 \leq i \leq n}$ of budgets is called admissible with respect to B and γ if $0 \leq x_1 \leq x_2 \leq \cdots \leq x_n = B$ and $p_1 \geq p_2 \geq \ldots \geq p_n$, where*

$$p_i = \mathsf{MP}\left(\mathsf{RT}\left(\mathcal{G}, \frac{x_i - x_{i-1}}{x_i - x_{i-1} + C_i - C_{i-1}}\right)\right) \quad (1)$$

for each $1 \leq i \leq n$, with $x_0 = 0$ and $C_0 = 0$. We denote by $\mathrm{ADM}(B,\gamma)$ the set of all admissible sequences with respect to B and γ.

We state our main result and the proof can be found in the full version.

Theorem 5 (Mean-payoff value of the partially-informed player). *Consider a strongly-connected first-price poorman mean-payoff bidding game $\mathcal{G}$. Let B be the initial budget of Max and γ be a finite budget distribution of Min with $\mathrm{supp}(\gamma) = \{C_1,\ldots,C_n\}$. Then*

$$\mathsf{MP}^{\downarrow}(\mathcal{G},\beta,\gamma) = \max_{(x_i)_{1 \leq i \leq n} \in \mathrm{ADM}(B,\gamma)} \mathsf{Val}(x_1,\ldots,x_n), \quad (2)$$

where $\mathsf{Val}(x_1,\ldots,x_n) = \sum_{i=1}^{n} \gamma(C_i) \cdot \mathsf{MP}\left(\mathsf{RT}\left(\mathcal{G}, \frac{x_i-x_{i-1}}{x_i-x_{i-1}+C_i-C_{i-1}}\right)\right)$ with $x_0 = 0$ and $C_0 = 0$.

We point to some interesting properties of Max's value:

Remark 2. Consider the bowtie game (Fig. 1) and assume Max's budget is fixed to $B = 1$ and Min's budget is drawn uniformly at random from $\{C_1, C_2\}$.

- When $C_1 = 1$ and $C_2 = 2$, the maximum is obtained at $x = 0.5$, thus Max's optimal expected payoff is $\frac{1}{3} = \frac{B}{B + C_2}$. We note that Max has a very simple optimal strategy in this case: "assume the worst" on Min's initial budget. That is, play according to an optimal strategy for initial budgets B and C_2.
- When $C_1 = 1$, and $C_2 = 5$, the maximum is obtained at $x = 0$. This is the dual of the case above. Max can "assume the best" on Min's initial budget and play according to an optimal strategy for budgets B and C_1. When Min's budget is C_1, this strategy guarantees a payoff of $\frac{B}{B + C_1}$. But when Min's budget is C_2, the strategy cannot guarantee a payoff above 0. Thus, the strategy guarantees an expected payoff of $\frac{1}{4} = \frac{1}{2} \cdot \frac{B}{B + C_1}$.
- There are cases in which Max's optimal strategy is not one of the trivial cases above. When $C_1 = 1$ and $C_2 = 3$, Max's optimal payoff is $\frac{1}{8}(5 - 2 \cdot \sqrt{2}) \approx 0.271$, which is strictly larger than both $\frac{1}{4} = \frac{1}{2} \cdot \frac{B}{B + C_1}$ and $\frac{1}{4} = \frac{B}{B + C_2}$. $\triangle$

All-pay poorman bidding We extend the technique in the previous section to all-pay poorman bidding. In order to state our results formally, we need to redefine the notion of admissible sequences since the optimal payoff that Max can guarantee under all-pay bidding differs from the payoff that he can guarantee under first-price bidding. Analogously to Def. 7 but now under all-pay bidding, we say that a sequence $(x_i)_{1 \leq i \leq n}$ of budgets is called *admissible* with respect to a budget B of Max and a budget distribution γ of Min if $0 \leq x_1 \leq x_2 \leq \cdots \leq x_n = B$ and $p_1 \geq p_2 \geq \ldots \geq p_n$, where now

$$p_i = \mathsf{MP}\Big(\mathsf{RT}\Big(\mathcal{G}, \Big(1 - \frac{C_i - C_{i-1}}{x_i - x_{i-1}}\Big) \cdot \mathbb{I}\Big(x_i - x_{i-1} > C_i - C_{i-1}\Big)\Big)\Big)$$

for each $1 \leq i \leq n$, with $x_0 = 0$ and $C_0 = 0$. Here, $\mathbb{I}$ is an indicator function that evaluates to 1 if the input logical formula is true, and to 0 if it is false. We are now ready to state our result on all-pay poorman mean-payoff bidding games.

Theorem 6 (Mean-payoff value of the partially-informed player)**.** *Consider a strongly-connected all-pay poorman mean-payoff bidding game $\mathcal{G}$. Let B be the initial budget of Max and γ be a finite budget distribution of Min with $\mathrm{supp}(\gamma) = \{C_1, \ldots, C_n\}$. Then*

$$\mathsf{MP}^{\downarrow}(\mathcal{G}, \beta, \gamma) = \max_{(x_i)_{1 \leq i \leq n} \in \mathrm{ADM}(B, \gamma)} \mathsf{Val}(x_1, \ldots, x_n), \quad (3)$$

where $\mathsf{Val}(x_1, \ldots, x_n) = \sum_{i=1}^{n} \gamma(C_i) \cdot \mathsf{MP}\Big(\mathsf{RT}\Big(\mathcal{G}, \Big(1 - \frac{C_i - C_{i-1}}{x_i - x_{i-1}}\Big) \cdot \mathbb{I}\Big(x_i - x_{i-1} > C_i - C_{i-1}\Big)\Big)\Big)$ *with $x_0 = 0$ and $C_0 = 0$ and $\mathbb{I}$ an indicator function.*

The Mean-Payoff Value of the Fully-Informed Player Under First-Price Poorman Bidding

In this section we identify the optimal expected payoff that the fully-informed player can guarantee in the bowtie game (Fig. 1) under first-price bidding. Suppose that Max's initial budget is B and Min's initial budget is drawn from a distribution γ. Consider the following collection of naive strategies for Min: when her initial budget is $C \in \mathrm{supp}(\gamma)$, Min plays according to an optimal full-information strategy for the ratio $\frac{B}{B + C}$.

We find it surprising that this collection of strategies is optimal for Min in the bowtie game. The technical challenge in this section is the lower bound. This result complements Thm. 5: we characterize both Min and Max's values in the bowtie game when the players are restricted to use pure strategies. We show, somewhat unexpectedly, that the two values do not necessarily coincide.

In order to state the result formally, we need the following definition. Intuitively, the *potential* of $\langle B, \gamma \rangle$ is the optimal expected payoff when Min plays according to the collection of naive strategies described above.

Definition 8. (Potential). *Given a budget $B \in \mathbb{R}$ of Max and a budget distribution γ with support $\mathrm{supp}(\gamma) = \{C_1, C_2, \ldots, C_k\}$ of Min, we define $\mathsf{Pot}(B, \gamma) = \sum_{j=1}^{k} \gamma(C_j) \cdot \frac{B}{B + C_j}$.*

The main result in this section is given in the following theorem, whose proof follows from Lemmas 8 and 9.

Theorem 7 (Mean-payoff value of the fully-informed player)**.** *Consider the bowtie game $\mathcal{G}_{\bowtie}$. Let B be the initial budget of Max and γ be a finite budget distribution of Min with $\mathrm{supp}(\gamma) = \{C_1, C_2, \ldots, C_k\}$. Then,*

$$\mathsf{MP}^{\uparrow}(\mathcal{G}, B, \gamma) = \mathsf{Pot}(B, \gamma) = \sum_{j=1}^{k} \gamma(C_j) \cdot \frac{B}{B + C_j}. \quad (4)$$

Before proving the theorem, we note the following.

Remark 3. (Inexistence of a value). Our result implies that the value in partial-information mean-payoff first-price poorman bidding games under pure strategies is not guaranteed to exist. Indeed, consider $\mathcal{G}_{\bowtie}$ with $B = 1$ and γ that draws Min's budget uniformly at random from $\{1, 2\}$. By Thm. 5, one can verify that the optimal choice of x is 1, thus $\mathsf{MP}^{\downarrow}(\mathcal{G}_{\bowtie}, B, \gamma) = \frac{1}{3}$. On the other hand, by Thm. 7, we have $\mathsf{MP}^{\uparrow}(\mathcal{G}_{\bowtie}, B, \gamma) = \frac{5}{12}$. $\triangle$

The upper bound is obtained when Min reveals her true budget immediately and plays according to the strategies described above. The following lemma follows from results on full-information games (Thm. 3).

Lemma 8 (Upper bound)**.** *For every $\epsilon > 0$, Min has a collection of strategies ensuring an expected payoff smaller than $\mathsf{Pot}(B, \gamma) + \epsilon$.*

We proceed to the more challenging lower bound and show that there are no Min strategies that perform better than the naive strategy above.

Lemma 9 (Lower bound). *For every $\epsilon > 0$ and for every collection $(g_j \in \mathcal{S}_{Min}(C_j))_{1 \le j \le k}$ of Min strategies, Max has a strategy ensuring an expected payoff greater than $\mathsf{Pot}(B, \gamma) - \epsilon$.*

Proof. Let $\epsilon > 0$, and let $(g_j \in \mathcal{S}_{Min}(C_j))_{1 \le j \le k}$ be a collection of Min strategies. We construct a counter strategy f of Max ensuring an expected payoff greater than $\mathsf{Pot}(B, \gamma) - \epsilon$. The proof is by induction over the size k of the support of γ. Obviously, if $k = 1$, Max has perfect information and can follow a full-information optimal strategy to guarantee a payoff of $\mathsf{Pot}(B, \gamma) = \frac{B}{B + C_1}$ (Thm. 3). So suppose that $k > 1$, and that the statement holds for every budget distribution of Min with a support strictly smaller than k.

Max carefully chooses a small part $x \le B$ of his budget and a part $y \le C_1$ of Min's budget. He plays according to a full-information strategy f for initial budgets x and y. This can result in three possible outcomes: (**O$_1$**) Min never uses more than y: the payoff is $\frac{x}{x+y}$ as in full-information games; (**O$_2$**) Min reveals her true initial budget, thus Max can distinguish between the case that Min's budget is C_i and C_j, and by the induction hypothesis he can ensure an expected payoff of $\mathsf{Pot}(B - x, \gamma)$ using his remaining budget; (**O$_3$**) Min does not reveal her true initial budget and spends more than y: Max's leftover budget is greater than $B - x$ and, for $1 \le j \le k$, when Min's budget is C_j, she has $C_j - y$, and Max re-starts the loop by selecting a new x.

We show that Max can choose x and y in a way that guarantees that the payoffs obtained in the first two outcomes are greater than the desired payoff $\mathsf{Pot}(B, \gamma) - \epsilon$. Also, when outcome O$_3$ occurs, the potential does not decrease, and eventually either O$_1$ or O$_2$ occur.

Formally, we describe a sequence $(\pi_i, B_i, \gamma_i)_{0 \le i \le m}$ of configurations comprising of a history π_i consistent with every strategy $(g_j)_{1 \le j \le k}$, the budget B_i of Max after π_i, and the budget distribution γ_i of Min with $\mathrm{supp}(\gamma_i) = \{C_1^i, C_2^i, \ldots, C_k^i\}$ following π_i. Tuple i represents the budget and budget distribution of the players following $i - 1$ choices of outcome O$_3$. Let $\lambda = 1 - \frac{\epsilon}{2}$ and $\rho = \frac{1}{\mathsf{Pot}(B,\gamma)} - 1$. We start with $(\pi_0, B_0, \gamma_0) = (v, B, \gamma)$ with v an initial vertex, and we show recursively how Max can update this tuple while ensuring that the following four properties are satisfied: (P$_1$) The history π_i is consistent with every $(g_j)_{1 \le j \le k}$; (P$_2$) Max spends his budget sufficiently slowly: $B_i \ge \lambda^i B$; (P$_3$) Min spends her budget sufficiently fast: $C_j^i \le C_j - \rho \cdot (1 - \lambda^i) B$ for every $1 \le j \le k$; (P$_4$) The potential never decreases: $\mathsf{Pot}(B_i, \gamma_i) \ge \mathsf{Pot}(B, \gamma)$.

Note that for the initial tuple $(\pi_0, B_0, \gamma_0) = (v, B, \gamma)$, these are trivially satisfied. Moreover, Property P$_3$ implies an upper bound on i, that is, outcome O$_3$ can happen only finitely many times: $\lim_{i \to \infty} C_1^i \le \lim_{i \to \infty} C_1 - \rho \cdot (1 - \lambda^i) B = C_1 - \rho \cdot B = B + C_1 - \frac{B}{\mathsf{Pot}(B,\gamma)} = \frac{1}{\frac{1}{B+C_1}} - \frac{1}{\sum_{j=1}^{k} \gamma(C_j) \cdot \frac{1}{B+C_j}}$ which is negative since $C_1 < C_2 < \ldots < C_k$, yet a negative C_1^i means that Min illegally bids higher than her available budget.

We now define the choices x_i and y_i for each $i \in \mathbb{N}$, and show that they satisfy the properties described above. Let

$x_i = \frac{\epsilon}{2} \cdot \lambda^i B$ and $y_i = \rho \cdot x_i$. For initial budgets x_i and y_i, let f_i be a Max strategy whose payoff is greater than $\frac{x_i}{x_i + y_i} - \epsilon$. Max follows f_i as long as Min spends at most y_i. Let $(\psi_j)_{1 \le j \le k}$ be plays such that for each $1 \le j \le k$ (1) the play $\pi_i \psi_j$ is consistent with the strategy g_j; (2) Max plays according to f_i along ψ_j; (3) ψ_j stops when Min spends more than y_i, and is infinite if she never does.

In the full version, we consider three possible cases, depending on whether the paths ψ_j are finite or infinite, and whether they are distinct or identical. If they are all infinite, or there are at least two distinct ones, we show that Max immediately has a way to obtain the desired payoff. If they are all identical and finite, we show that, while Max cannot immediately get the desired payoff, he can go to the next step by setting $\phi_{i+1} = \phi_i \psi_1$, and restarting. $\square$

Discussion and Future Work

We initiate the study of partial-information bidding games, by studying games with partially-observed budgets. For mean-payoff games, we show a complete picture in strongly-connected games for the partially-informed player, which is the more important case in practice. By identifying the value for the fully-informed player in the bowtie game, we show that the value in mean-payoff bidding games does not necessarily exist when restricting to pure strategies.

We discuss open problems in this model. First, we focus on games played on strongly-connected graphs. Reasoning about such games is the crux of the solution to general full-information bidding games. We thus expect that our results will be key in the solution of partial-information bidding games on general graphs. This extension, however, is not straightforward as in the full-information setting, and we leave it as an open question. Second, we identify the value of the fully-informed player in the bowtie game $\mathcal{G}_{\bowtie}$. Reasoning about $\mathcal{G}_{\bowtie}$ was the crux of the solution to general strongly-connected full-information bidding games. In fact, the same technique was used to lift a solution for $\mathcal{G}_{\bowtie}$ to general strongly-connected games under all the previously-studied bidding mechanisms. In partial-information games, however, this technique breaks the intricate analysis in the proof of Thm. 7. Again, we expect a solution to the bowtie game to be a key ingredient in the solution to general strongly-connected games, and we leave the problem open. Finally, we showed that the value does not necessarily exist under pure strategies. We leave open the problem of developing optimal mixed strategies for the players.

This work is part of a research that combines formal methods and AI including multi-agent graph games (Alur, Henzinger, and Kupferman 2002), logics to reason about strategies (Chatterjee, Henzinger, and Piterman 2010; Mogavero et al. 2014) and in particular, their application in auctions (Mittelmann et al. 2022), enhancing *network-formation games* with concepts from formal methods (e.g., (Avni, Kupferman, and Tamir 2016)), and many more.

References

Aghajohari, M.; Avni, G.; and Henzinger, T. A. 2021. Determinacy in Discrete-Bidding Infinite-Duration Games. *Log. Methods Comput. Sci.*, 17(1).

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

Aumann, R. J.; Maschler, M.; and Stearns, R. E. 1995. *Repeated games with incomplete information*. MIT press.

Avni, G.; and Henzinger, T. A. 2020. A Survey of Bidding Games on Graphs. In *Proc. 31st CONCUR*, volume 171 of *LIPIcs*, 2:1–2:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

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

Avni, G.; Henzinger, T. A.; and Ibsen-Jensen, R. 2018. Infinite-Duration Poorman-Bidding Games. In *Proc. 14th WINE*, volume 11316 of *LNCS*, 21–36. Springer.

Avni, G.; Henzinger, T. A.; and Žikelić, Đ. 2021. Bidding mechanisms in graph games. *J. Comput. Syst. Sci.*, 119: 133–144.

Avni, G.; Ibsen-Jensen, R.; and Tkadlec, J. 2020. All-Pay Bidding Games on Graphs. In *Proc. 34th AAAI*, 1798–1805. AAAI Press.

Avni, G.; Jecker, I.; and Žikelić, Đ. 2021. Infinite-Duration All-Pay Bidding Games. In *Proc. 32nd SODA*, 617–636.

Avni, G.; Kupferman, O.; and Tamir, T. 2016. Network-formation games with regular objectives. *Inf. Comput.*, 251: 165–178.

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

Chatterjee, K.; Henzinger, T. A.; and Piterman, N. 2010. Strategy logic. *Inf. Comput.*, 208(6): 677–693.

Condon, A. 1992. The Complexity of Stochastic Games. *Inf. Comput.*, 96(2): 203–224.

Develin, M.; and Payne, S. 2010. Discrete bidding games. *Electron. J. Comb.*, 17(1).

Lazarus, A. J.; Loeb, D. E.; Propp, J. G.; Stromquist, W. R.; and Ullman, D. H. 1999. Combinatorial Games under Auction Play. *Games and Economic Behavior*, 27: 229–264.

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

Meir, R.; Kalai, G.; and Tennenholtz, M. 2018. Bidding games and efficient allocations. *Games Econ. Behav.*, 112: 166–193.

Mittelmann, M.; Maubert, B.; Murano, A.; and Perrussel, L. 2022. Automated Synthesis of Mechanisms. In *Proc. 31st IJCAI*, 426–432. ijcai.org.

Mogavero, F.; Murano, A.; Perelli, G.; and Vardi, M. Y. 2014. Reasoning About Strategies: On the Model-Checking Problem. *ACM Trans. Comput. Log.*, 15(4): 34:1–34:47.

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

Pnueli, A.; and Rosner, R. 1989. On the Synthesis of a Reactive Module. In *Proc. 16th POPL*, 179–190. ACM Press.

Puterman, M. L. 2005. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY, USA: John Wiley & Sons, Inc. ISBN 978-0-471-72782-8.

Raskin, J.; Chatterjee, K.; Doyen, L.; and Henzinger, T. A. 2007. Algorithms for Omega-Regular Games with Imperfect Information. *Log. Methods Comput. Sci.*, 3(3).

Reif, J. H. 1984. The Complexity of Two-Player Games of Incomplete Information. *J. Comput. Syst. Sci.*, 29(2): 274–301.

Wulf, M. D.; Doyen, L.; and Raskin, J. 2006. A Lattice Theory for Solving Games of Imperfect Information. In *Proc. 9th HSCC*, volume 3927 of *LNCS*, 153–168. Springer.