Bidding Mechanisms in Graph Games

Guy Avni
IST Austria, Klosterneuburg, Austria
guy.avni@ist.ac.at

Thomas A. Henzinger
IST Austria, Klosterneuburg, Austria
tah@ist.ac.at

Đorđe Žikelić
IST Austria, Klosterneuburg, Austria
djordje.zikelic@ist.ac.at

Abstract

In two-player games on graphs, the players move a token through a graph to produce a finite or infinite path, which determines the qualitative winner or quantitative payoff of the game. We study *bidding games* in which the players bid for the right to move the token. Several bidding rules were studied previously. In *Richman* bidding, in each round, the players simultaneously submit bids, and the higher bidder moves the token and pays the other player. *Poorman* bidding is similar except that the winner of the bidding pays the "bank" rather than the other player. *Taxman* bidding spans the spectrum between Richman and poorman bidding. They are parameterized by a constant $\tau \in [0, 1]$: portion τ of the winning bid is paid to the other player, and portion $1 - \tau$ to the bank. While finite-duration (reachability) taxman games have been studied before, we present, for the first time, results on *infinite-duration* taxman games. It was previously shown that both Richman and poorman infinite-duration games with qualitative objectives reduce to reachability games, and we show a similar result here. Our most interesting results concern quantitative taxman games, namely *mean-payoff* games, where poorman and Richman bidding differ significantly. A central quantity in these games is the *ratio* between the two players' initial budgets. While in poorman mean-payoff games, the optimal payoff of a player depends on the initial ratio, in Richman bidding, the payoff depends only on the structure of the game. In both games the optimal payoffs can be found using (different) probabilistic connections with *random-turn* games in which in each turn, instead of bidding, a coin is tossed to determine which player moves. While the value with Richman bidding equals the value of a random-turn game with an un-biased coin, with poorman bidding, the bias in the coin is the initial ratio of the budgets. We give a complete classification of mean-payoff taxman games that is based on a probabilistic connection: the value of a taxman bidding game with parameter τ and initial ratio r, equals the value of a random-turn game that uses a coin with bias $F(\tau, r) = \frac{r + \tau \cdot (1-r)}{1 + \tau}$. Thus, we show that Richman bidding is the exception; namely, for every $\tau < 1$, the value of the game depends on the initial ratio. Our proof technique simplifies and unifies the previous proof techniques for both Richman and poorman bidding.

1   Introduction

Two-player infinite-duration games on graphs are a central class of games in formal verification [2], where they are used, for example, to solve synthesis [19], and they have deep connections to foundations of logic [21]. A graph game proceeds by placing a token on a vertex in the graph, which the players move throughout the graph to produce an infinite path ("play") π. The game is zero-sum and π determines the winner or payoff. Graph games can be classified according to the players' objectives. For example, the simplest objective is *reachability*, where Player 1 wins iff an infinite path visits a designated target vertex. Another classification of graph games is the *mode of moving* the token. The most studied mode of moving is *turn based*, where the players alternate turns in moving the token.

In *bidding games*, in each turn, an "auction" is held between the two players in order to determine which player moves the token. The bidding mode of moving was introduced in [13, 14] for reachability games, where the following bidding rules where defined. In *Richman* bidding (named after David Richman), each player has a budget, and before each turn, the players submit bids simultaneously, where a bid is legal if it does not exceed the available budget. The player who bids higher wins the bidding, pays the bid to the other player, and moves the token. A second bidding rule called *poorman* bidding in [13], is similar except that the winner of the bidding pays the "bank" rather than the other player. Thus, the bid is deducted from his budget and the money is lost. A third bidding rule on which we focus in this paper, called *taxman* in [13] spans the spectrum between poorman and Richman bidding. Taxman bidding is parameterized by $\tau \in [0, 1]$: the winner of a bidding pays portion τ of his bid to the other player and portion $1 - \tau$ to the bank. Taxman bidding with $\tau = 1$ coincides with Richman bidding and taxman bidding with $\tau = 0$ coincides with poorman bidding.

Bidding games are relevant for several communities in Computer Science. In formal methods, graph games are used to reason about systems. Poorman bidding games naturally model concurrent systems where processes pay the scheduler for moving. Block-chain technology like Etherium is an example of such a system, which is a challenging to formally verify [9, 3]. In Algorithmic Game Theory [17], auction design is a central research topic that is motivated by the abundance of auctions for online advertisements [16]. Infinite-duration bidding games can model ongoing auctions and can be used to devise bidding strategies for objectives like: "In the long run, an advertiser's ad should show at least half of the time". In Artificial Intelligence, bidding games with Richman bidding have been used to reason about combinatorial negotiations [15]. Finally, *discrete-bidding* games [11], in which the granularity of the bids is restricted by assuming that the budgets are given using coins, have been studied mostly for recreational games, like bidding chess [6].

Both Richman and poorman infinite-duration games have a surprising, elegant, though different, mathematical structure as we elaborate below. Our study of taxman bidding aims at a better understanding of this structure and at shedding light on the differences between the seemingly similar bidding rules.

A central quantity in bidding games is the *initial ratio* of the players budgets. Formally, assuming that, for $i \in \{1, 2\}$, Player i's initial budget is B_i, we say that Player 1's initial ratio is $B_1/(B_1 + B_2)$. The central question that was studied in [13] regards the existence of a necessary and sufficient initial ratio to guarantee winning the game. Formally, the *threshold ratio* in a vertex v, denoted $\text{Th}(v)$, is such that if Player 1's initial ratio exceeds $\text{Th}(v)$, he can guarantee winning the game, and if his initial ratio is less than $\text{Th}(v)$, Player 2 can guarantee

winning the game[1]. Existence of threshold ratios in reachability games for all three bidding mechanisms was shown in [13].

Reachability Richman-bidding games have an interesting probabilistic connection [14]. To state the connection, we first need to introduce *random-turn* games. Let $p \in [0, 1]$. In a random-turn game that is parameterized by p, in each turn, rather than bidding, the player who moves is chosen by throwing a (possibly) biased coin: with probability p, Player 1 chooses how to move the token, and Player 2 chooses with probability $1 - p$. Formally, a random-turn game is a special case of a stochastic game [10]. Consider a reachability Richman-bidding game $\mathcal{G}$. We construct a "uniform" random-turn game on top of $\mathcal{G}$, denoted $\mathtt{RT}^{0.5}(\mathcal{G})$, in which we throw an unbiased coin in each turn. The objective of Player 1 remains reaching his target vertex. It is well known that each vertex in $\mathtt{RT}^{0.5}(\mathcal{G})$ has a *value*, which is, informally, the probability of reaching the target when both players play optimally, and which we denote by $val(\mathtt{RT}^{0.5}(\mathcal{G}), v)$. We are ready to state the probabilistic connection: For every vertex v in the Richman game $\mathcal{G}$, the threshold ratio in v equals $1 - val(\mathtt{RT}^{0.5}(\mathcal{G}), v)$. We note that such a connection is not known and is unlikely to exist in reachability games with neither poorman nor taxman bidding. Random-turn games have been extensively studied in their own right, mostly with unbiased coin tosses, since the seminal paper [18].

Infinite-duration bidding games have been recently studied with Richman [4] and poorman [5] bidding. For qualitative objectives, namely games in which one player wins and the other player loses, both bidding rules have similar properties. By reducing general qualitative games to reachability games, it is shown that threshold ratios exist for both types of bidding rules. We show a similar result for qualitative games with taxman bidding.

Things get interesting in *mean-payoff* games, which are quantitative games: an infinite play has a *payoff*, which is Player 1's reward and Player 2's cost (see an example of a mean-payoff game in Figure 1). We thus call the players in a mean-payoff game Max and Min, respectively. We focus on games that are played on strongly-connected graphs. With Richman bidding [4], the initial budget of the players does not matter: A mean-payoff Richman-bidding game $\mathcal{G}$ has a value $c \in \mathbb{R}$ that depends only on the structure of the game such that Min can guarantee a cost of at most c with any positive budget, and with any positive budget, Max can guarantee a payoff of at least $c - \epsilon$, for every $\epsilon > 0$. Moreover, the value c of $\mathcal{G}$ equals the value of a random-turn game $\mathtt{RT}^{0.5}(\mathcal{G})$ that is constructed on top of $\mathcal{G}$. Since $\mathcal{G}$ is a mean-payoff game, $\mathtt{RT}^{0.5}(\mathcal{G})$ is a mean-payoff stochastic game, and its value, which again, is a well-known concept, is the expected payoff when both players play optimally.

Mean-payoff poorman-bidding games have different properties. Unlike with Richman bidding, the value of the game depends on the initial ratio. That is, with a higher initial ratio, Max can guarantee a better payoff. While the probabilistic connection for mean-payoff Richman games is not entirely unexpected given the probabilistic connection for reachability Richman games, we find it surprising that mean-payoff poorman games exhibit a probabilistic connection, which is in fact richer than for Richman bidding. The connection for poorman games is the following: Suppose Max's initial ratio is $r \in [0, 1]$ in a game $\mathcal{G}$. Then, the value in $\mathcal{G}$ with respect to r is the value of the random-turn game $\mathtt{RT}^r(\mathcal{G})$ in which in each turn, we toss a biased coin that chooses Max with probability r and Min with probability $1 - r$.

---

[1] When the initial ratio is exactly $\mathtt{Th}(v)$, the winner depends on the mechanism with which ties are broken. Our results do not depend on a specific tie-breaking mechanism.Tie-breaking mechanisms are particularly important in discrete-bidding games [1].
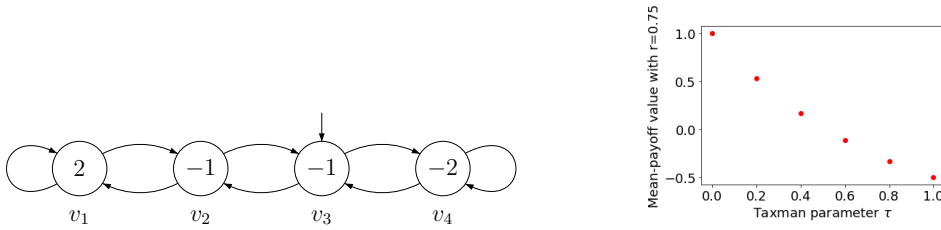
Figure 1 On the left, a mean-payoff game $\mathcal{G}$. On the right, the mean-payoff value of $\mathcal{G}$, where the initial ratio is fixed to 0.75 and the taxman parameter τ varies. The value of $\mathcal{G}$ with Richman bidding is -0.5, with poorman bidding, it is 1, and, for example, with $\tau = 0.2$, it is 0.533.

Given this difference between the two bidding rules, one may wonder how do mean-payoff taxman games behave, since these bidding rules span the spectrum between Richman and poorman bidding. Our main contribution is a complete solution to this question: we identify a probabilistic connection for a taxman game $\mathcal{G}$ that depends on the parameter τ of the bidding and the initial ratio r. That is, we show that the value of the game equals the value of the random-turn game $\mathtt{RT}^{F(\tau,r)}(\mathcal{G})$, where $F(\tau, r) = \frac{r + \tau \cdot (1-r)}{1+\tau}$. The construction gives rise to optimal strategies w.r.t. τ and the initial ratio. As a sanity check, note that for $\tau = 1$, we have $F(\tau, r) = 0.5$, which agrees with the result on Richman bidding, and for $\tau = 0$, we have $F(\tau, r) = r$, which agrees with the result on poorman bidding. In Figure 1, we depict some mean-payoff values for a fixed initial ratio and varying taxman parameter. Previous results only give the two endpoints in the plot, and the mid points in the plot are obtained using the results in this paper.

The main technical challenge is constructing an optimal strategy for Max, which, intuitively, performs a de-randomization; with a deterministic bidding strategy, Max guarantees that the ratio of the time that is spent in each vertex is the same as in a random behavior. The construction of Max strategy involves two components. First, we assign an "importance" to each vertex v, which we call *strength* and denote $\mathrm{St}(v)$. Intuitively, if $\mathrm{St}(v) > \mathrm{St}(u)$, then it is more important for Max to move in v than in u. Second, when the game reaches a vertex v, Max's bid is a careful normalization of $\mathrm{St}(v)$ so that changes in Max's ratio are matched with the accumulated weights in the game. Finding the right normalization is intricate and it consists of the main technical contribution of this paper. Previous such normalizations were constructed for Richman and poorman mean-payoff games [4, 5]. The construction for Richman bidding is much more complicated than the one we present here. The construction for poorman bidding is ad-hoc and does not generalize. Our construction for taxman bidding thus unifies these constructions and simplifies them. It uses techniques that can generalize beyond taxman bidding. Finally, we study, for the first time, complexity problems for taxman games.

Due to lack of space, some proofs appear in the full version.

2   Preliminaries

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

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

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

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

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

Strategies and plays.   A *strategy* is a recipe for how to play a game. It is a function that, given a finite *history* of the game, prescribes to a player which *action* to take, where we define these two notions below. For example, in turn-based games, a strategy takes as input, the sequence of vertices that were visited so far, and it outputs the next vertex to move to. In bidding games, histories and strategies are more involved as they maintain the information about the bids and winners of the bids. Formally, a history in a bidding game is $\pi = \langle v_1, b_1, i_1 \rangle, \ldots, \langle v_k, b_k, i_k \rangle, v_{k+1} \in (V \times \mathbb{R} \times \{1, 2\})^* \cdot V$, where for $1 \leq j \leq k+1$, the token is placed on vertex v_j at round j, for $1 \leq j \leq k$, the winning bid is b_j and the winner is Player i_j. Consider a finite history π. For $i \in \{1, 2\}$, let $W_i(\pi) \subseteq \{1, \ldots, k\}$ denote the indices in which Player i is the winner of the bidding in π. Let B_i^I be the initial budget of Player i. Player i's budget following π, denoted $B_i(\pi)$, depends on the bidding mechanism. For example, in Richman bidding, $B_1(\pi) = B_i^I - \sum_{j \in W_1(\pi)} b_j + \sum_{j \in W_2(\pi)} b_j$, B_2 is defined dually, and the definition is similar for taxman and poorman bidding. Given a history π that ends in v, a strategy for Player i prescribes an action $\langle b, v \rangle$, where $b \leq B_i(\pi)$ is a bid that does not exceed the available budget and v is a vertex to move to upon winning, where we require that v is a neighbor of v_{k+1}. An initial vertex, initial budgets, and two strategies for the players determine a unique infinite *play* π for the game. The vertices that π visits form an infinite path $path(\pi)$.

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

1. In *reachability games*, Player 1 has a target vertex t and an infinite play is winning iff it visits t.

2. In *parity games*, each vertex is labeled with an index in $\{1, \ldots, d\}$. An infinite path is winning for Player 1 iff the parity of the maximal index that is visited infinitely often is odd.

3. *Mean-payoff games* are played on weighted directed graphs, with weights given by a function $w : V \to \mathbb{Q}$. Consider an infinite path $\eta = v_1, v_2, \cdots \in V^\omega$. For $n \in \mathbb{N}$, the prefix of length n of η is η^n, and we define its *energy* to be $E(\eta^n) = \sum_{i=1}^{n} w(v_i)$. The *payoff* of η is $\mathrm{MP}(\eta) = \liminf_{n \to \infty} E(\eta^n)/n$. Player 1 wins η iff $\mathrm{MP}(\eta) \geq 0$.

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

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

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

Random-turn games. A *stochastic game* [10] is a graph game in which the vertices are partitioned between two players and a *nature* player. As in turn-based games, whenever the game reaches a vertex that is controlled by Player i, for $i = 1, 2$, he choses how the game proceeds, and whenever the game reaches a vertex v that is controlled by nature, the next vertex is chosen according to a probability distribution that depends only on v.

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

The following definitions are standard, and we refer the reader to [20] for more details. A strategy in a stochastic game is similar to a turn-based game; namely, given the history of vertices visited so far, the strategy choses the next vertex. Fixing two such strategies f and g for both players gives rise to a distribution $D(f, g)$ on infinite paths. Intuitively, Player 1's goal is to maximize the probability that his objective is met. An *optimal* strategy for Player 1 guarantees that the objective is met with probability at least c and, intuitively, he cannot do better, thus Player 2 has a strategy that guarantees that the objective is violated with probability at least $(1 - c)$. It is well known that optimal positional strategies exist for the objectives that we consider.

▶ **Definition 3** (Values in stochastic games). *Consider a bidding game $\mathcal{G}$, let $r \in [0, 1]$, and consider two optimal strategies f and g for the two players in $\mathrm{RT}^r(\mathcal{G})$. When $\mathcal{G}$ is a qualitative game with objective O, the* value *of $\mathrm{RT}^r(\mathcal{G})$, denoted $val(\mathrm{RT}^r(\mathcal{G}))$, is $\mathrm{Pr}_{\eta \sim D(f,g)} \mathrm{Pr}[\eta \in O]$. When $\mathcal{G}$ is a mean-payoff game, the* mean-payoff value *of $\mathrm{RT}^r(\mathcal{G})$, denoted $MP(\mathrm{RT}^r(\mathcal{G}))$, is $\mathbb{E}_{\eta \in D(f,g)} MP(\eta)$.*

3   Qualitative Taxman Games

In [14, 13], reachability bidding games were studied with a slightly different definition, which we call *double-reachability*: both players have a target, where we denote by t_i the target of Player i, for $i \in \{1, 2\}$, all vertices have a path to both targets, and the game ends once one of the targets is reached. They show the following results.

▶ **Theorem 4.** *[14, 13] Consider a double-reachability bidding game $\mathcal{G}$ and a vertex v. The threshold ratio exists in v with Richman, poorman, and taxman bidding. Moreover, threshold ratios have the following properties. For the target vertex t_1 of Player 1, we have $Th(t_1) = 0$, and for the target t_2 of Player 2, we have $Th(v) = 1$. Consider some other vertex v and denote $v^+, v^- \in N(v)$ the vertices with the minimal and maximal thresholds in the neighborhood of v, thus for every $u \in N(v)$, we have $Th(v^-) \leq Th(u) \leq Th(v^+)$.*

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

Moreover, only double-reachability Richman-bidding games exhibit the following probabilistic connection: for every vertex v, we have $Th(v) = 1 - val(RT^{0.5}(\mathcal{G}), v)$. Thus, for games played on finite graphs, the threshold ratios are all rational numbers. However, threshold ratios with poorman-bidding need not be rational in finite games.

The equivalence between double-reachability bidding games and reachability games with Richman- and poorman-bidding is shown in [4] and [5]. The following proposition is the key component in showing the equivalence as well as in the reduction from parity taxman games to reachability taxman games.

▶ **Lemma 5.** *Consider a reachability taxman game $\mathcal{G}$. Suppose that every vertex in $\mathcal{G}$ has a path to the target of Player 1. Then, for any taxman parameter, Player 1 wins from every vertex with any positive initial budget. Thus, for every vertex v, we have $Th(v) = 0$.*

Proof. Let $\mathcal{G} = \langle V, E, t \rangle$, where $n = |V| - 1$. Suppose the game starts from a vertex v, and let $\epsilon > 0$ be the initial budget of Player 1. Since there is a path from v to Player 1's target, there is a path of length at most n. Thus, if Player 1 wins n consecutive biddings, he wins the game. Intuitively, Player 1 carefully chooses n increasing bids such that if Player 2 wins one of these bids, Player 1's ratio increases by a constant over his initial budget. By repeatedly playing according to such a strategy, Player 1 guarantees that his ratio increases and will eventually allow him to win n biddings in a row. Formally, if $\tau = 0$, then $\mathcal{G}$ is a Richman game and the proof of the lemma can be found in [4]. Otherwise, pick a sufficiently large $r \in \mathbb{N}$ such that $\tau > \frac{2}{r-1}$ and $r \geq 3$. Fix $0 < m < \frac{\epsilon}{r^n}$. Player 1 proceeds as follows: after winning i times, for $0 \leq i$, he bids $m \cdot r^i$ and, upon winning the bidding, he moves towards t along any shortest path. Since $m + mr + \cdots + mr^{n-1} < mr^n < \epsilon$, Player 1 has sufficient budget to win n consecutive biddings. If Player 2 does not win any of the first n biddings, Player 1 wins the game. On the other hand, if Player 2 wins the k-th bidding with $1 \leq k \leq n$, we show in the full version that his ratio increases by a fixed amount $b = \frac{mr}{(1-\epsilon)(r-1)} > 0$. ◄

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

▶ **Corollary 6.** *Consider a reachability taxman game $\mathcal{G} = \langle V, E, t \rangle$. Let $S \subseteq V$ be the set of vertices that have no path to t. Let $T \subseteq V$ be a set of vertices such that $t \in T$ and every*

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

The following theorem, whose proof can be found in the full version, uses Lemma 5 to classify the bottom-strongly-connected components of a parity taxman game as those that are winning and losing for Player 1, thereby constructing a reachability taxman game.

▶ **Theorem 7.** *Parity taxman games are linearly reducible to reachability taxman games. Specifically, threshold ratios exist in parity taxman games.*

4   Mean-Payoff Taxman Games

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

4.1   Strongly-connected mean-payoff taxman games

We start by formally defining the value of a strongly-connected mean-payoff game. Lemma 5 implies that in a strongly-connected game, a player can draw the game from every vertex to any other vertex with any positive initial budget. Since mean-payoff objectives are prefix independent, it follows that the vertex from which the game starts does not matter. Indeed, if the game starts at a vertex v with Max having initial ratio $r + \epsilon$, then Max can use $\epsilon/2$ of his budget to draw the game to a vertex u and continue as if he starts the game with initial ratio $r + \epsilon/2$.

▶ **Definition 8** (Mean-payoff value)**.** *Consider a strongly-connected mean-payoff game $\mathcal{G}$, a ratio $r \in (0, 1)$, and a taxman parameter $\tau \in [0, 1]$. The mean-payoff value of $\mathcal{G}$ w.r.t. r and τ, is a value $c \in \mathbb{R}$ such that for every $\epsilon > 0$*

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

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

▶ **Theorem 9.** *Consider a strongly-connected mean-payoff taxman game $\mathcal{G}$ with taxman parameter $\tau \in [0, 1]$ and an initial ratio $r \in (0, 1)$. The value of $\mathcal{G}$ w.r.t. τ and r equals the value of the random-turn game $\mathit{RT}^{F(\tau,r)}(\mathcal{G})$ in which Max is chosen to move with probability $F(\tau, r)$ and Min with probability $1 - F(\tau, r)$, where $F(\tau, r) = \frac{r + \tau(1-r)}{1+\tau}$.*

We show that in order to prove Theorem 9, it suffices to prove the following intermediate lemma, whose proof can be found in the full version, and follows from the advantage of Min in the definition of payoff.

▶ **Lemma 10.** *Consider a strongly-connected mean-payoff taxman game $\mathcal{G}$, a taxman parameter τ, and an initial ratio $r \in (0,1)$ such that $\text{MP}(RT^{F(\tau,r)}) = 0$ for $F(\tau,r) = \frac{r+\tau(1-r)}{1+\tau}$. Then, for every $\epsilon > 0$ Max has a strategy that guarantees that no matter how Min plays, the payoff is greater than $-\epsilon$.*

4.2 The importance of moving

The first part of the construction of an optimal strategy for Max as in Lemma 10 is to assign, to each vertex $v \in V$, a *strength*, denoted $\text{St}(v)$, where $\text{St}(v) \in \mathbb{Q}_{\geq 0}$. Intuitively, if $\text{St}(v) > \text{St}(u)$, for $u, v \in V$, it is more important for Max to move in v than it is in u. We follow the construction in [5], which uses the concept of *potentials*, which is a well-known concept in stochastic games (see [20]) and was originally defined in the context of the strategy iteration algorithm [12]. For completeness, we present the definitions below.

Consider a strongly-connected mean-payoff game $\mathcal{G}$, and let $p \in [0,1]$. Let f and g be two optimal positional strategies in $RT^p(\mathcal{G})$, for Min and Max, respectively. For a vertex $v \in V$, let $v^-, v^+ \in V$ be such that Max proceeds from v to v^+ according to g and Min proceeds from v to v^- according to f. It is not hard to see that the mean-payoff value in all vertices in $RT^p(\mathcal{G})$ is the same and we denote it by $\text{MP}(RT^p(\mathcal{G}))$. We denote the potential of v by $\text{Pot}^p(v)$ and the strength of v by $\text{St}^p(v)$, and we define them as follows.

$$\text{Pot}^p(v) = p \cdot \text{Pot}^p(v^+) + (1-p) \cdot \text{Pot}^p(v^-) + w(v) - \text{MP}(RT^p(\mathcal{G})) \text{ and}$$
$$\text{St}^p(v) = p \cdot (1-p) \cdot \left(\text{Pot}^p(v^+) - \text{Pot}^p(v^-)\right)$$

There are optimal strategies for which $\text{Pot}^p(v^-) \leq \text{Pot}^p(v') \leq \text{Pot}^p(v^+)$, for every $v' \in N(v)$, which can be found, for example, using the strategy iteration algorithm. Note that $\text{St}(v) \geq 0$, for every $v \in V$.

Consider a finite path $\pi = v_1, \ldots, v_n$ in $\mathcal{G}$. We intuitively think of π as a play, where for every $1 \leq i < n$, the bid of Max in v_i is $\text{St}(v_i)$ and he moves to v_i^+ upon winning. Thus, if $v_{i+1} = v_i^+$, we say that Max won in v_i, and if $v_{i+1} \neq v_i^+$, we say that Max lost in v_i. Let $W(\pi)$ and $L(\pi)$ respectively be the indices in which Max wins and loses in π. We call Max wins *investments* and Max loses *gains*, where intuitively he *invests* in increasing the energy and *gains* a higher ratio of the budget whenever the energy decreases. Let $G(\pi)$ and $I(\pi)$ be the sum of gains and investments in π, respectively, thus $G(\pi) = \sum_{i \in L(\pi)} \text{St}(v_i)$ and $I(\pi) = \sum_{i \in W(\pi)} \text{St}(v_i)$. Recall that the energy of π is $E(\pi) = \sum_{1 \leq i < n} w(v_i)$. The following lemma, whose proof can be found in the full version, which generalizes a similar lemma in [5], connects the strength with the change in energy.

▶ **Lemma 11.** *Consider a strongly-connected mean-payoff game $\mathcal{G}$ and $p \in [0,1]$. For every finite path $\pi = v_1, \ldots, v_n$ in $\mathcal{G}$, we have $\text{Pot}^p(v_1) - \text{Pot}^p(v_n) + (n-1) \cdot \text{MP}(RT^p(\mathcal{G})) \leq E(\pi) + G(\pi)/(1-p) - I(\pi)/p$. In particular, when $p = \nu/(\mu+\nu)$ for $\nu, \mu > 0$, there is a constant $P = \min_v \text{Pot}^p(v) - \max_v \text{Pot}^p(v)$ such that $\frac{\nu \cdot \mu}{\nu+\mu} \cdot \left(E(\pi) - P - (n-1) \cdot \text{MP}(RT^{\frac{\nu}{\mu+\nu}}(\mathcal{G}))\right) \geq \mu \cdot I(\pi) - \nu \cdot G(\pi)$.*

4.3 Normalizing the bids

Whenever the game reaches a vertex v, Max's bid is obtained by carefully normalizing the strength of v. More formally, assuming an initial ratio r, in v, Max bids $r \cdot (1-r) \cdot \text{St}(v) \cdot \beta_x$, where β_x is the normalization factor and $x \in \mathbb{R}_{\geq 1}$. In this section we show how to choose the normalization factor. We associate with every $x \geq 1$, two numbers: a ratio r_x and β_x both in $(0,1)$. We think of $(r_x)_{x \geq 1}$ as a sequence and a play gives rise to a walk on the

sequence, which corresponds to the changes in energy in the bidding game. When the walk is in $x \geq 1$, Max uses the normalization factor β_x. If Max wins a bidding, we take a step up on the sequence, modeling the increase of energy, and when Min wins, we talk a step down. The size of the step depends on the strength of v. We prove existence of sequences with properties given in the following lemma and formally define Max's strategy after it.

▶ **Lemma 12.** *Consider a game $\mathcal{G}$, a finite set of non-negative strengths $S \subseteq \mathbb{R}_{\geq 0}$, a ratio $r \in (0,1)$, and a taxman parameter $\tau \in [0,1]$. For every $K > \frac{\tau r^2 + r(1-r)}{\tau(1-r)^2 + r(1-r)}$ there exist sequences $(r_x)_{x \geq 1}$ and $(\beta_x)_{x \geq 1}$ with the following properties.*

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

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

We first show how Lemma 12 implies Theorem 9.

Proof that Lemma 12 implies Lemma 10. Fix $\epsilon > 0$, we construct strategy for Max guaranteeing a payoff greater than $-\epsilon$, as wanted. Observe that

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

Thus, since by assumption $\mathrm{MP}(\mathrm{RT}^{F(\tau,r)}(\mathcal{G})) = 0$ and $\mathrm{MP}(\mathrm{RT}^p(\mathcal{G}))$ is a continuous function in $p \in [0, 1]$ [8, 22], we can pick $K > F(\tau, r)$ such that $\mathrm{MP}(\mathrm{RT}^{\frac{r}{r + (1-r)K}}(\mathcal{G})) > -\epsilon$.

We now describe Max's strategy. We think of the change in Max's ratio as a walk on $\mathbb{R}_{\geq 1}$. Each position $x \in \mathbb{R}_{\geq 1}$ is associated with a ratio r_x. The walk starts in a position x_0 such that Max's initial ratio is at least r_{x_0}. Let $\nu = r$ and $\mu = K(1 - r)$. Suppose the token is placed on a vertex $v \in V$. Then, Max's bid is $r \cdot (1 - r) \cdot \beta_x \cdot \mathrm{St}(v)$, where the ratios of Max and Min are normalized to sum up to 1, and he proceeds to v^+ upon winning. If Max wins, the walk proceeds up $\mu \cdot \mathrm{St}(v)$ steps to $x + \mu \mathrm{St}(v)$, and if he loses, the walk proceeds down to $x - \nu \mathrm{St}(v)$. Suppose Min fixes some strategy and let $\pi = v_1, \ldots, v_n$ be a finite prefix of the play that is generated by the two strategies. Suppose the walk following π reaches $x \in \mathbb{R}$. Then, using the terminology of the previous section, we have $x = x_0 - G(\pi) \cdot \nu + I(\pi) \cdot \mu$. Lemma 12 shows that the walk always stays above 1, thus $x \geq 1$. Combining with Lemma 11, we get $\frac{\nu + \mu}{\nu \cdot \mu}(1 - x_0) + P + (n - 1) \cdot \mathrm{MP}(\mathrm{RT}^{\frac{\nu}{\nu + \mu}}(\mathcal{G})) \leq E(\pi)$. Thus, dividing both sides by n and letting $n \to \infty$, since x_0 and P are constants depending only on K we conclude that this strategy guarantees payoff at least $\mathrm{MP}(\mathrm{RT}^{\frac{\nu}{\nu + \mu}}(\mathcal{G})) > -\epsilon$. ◀

We continue to prove Lemma 12.

Proof of Lemma 12. Note that $\frac{\tau r^2 + r(1-r)}{\tau(1-r)^2 + r(1-r)}$ is well-defined for $r \in (0, 1)$. Fix $\tau \in [0, 1]$ and $r \in (0, 1)$. Let $K > \frac{\tau r^2 + r(1-r)}{\tau(1-r)^2 + r(1-r)}$. Observe that the two inequalities in Point 4 are equivalent to:

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

Point 3 combined with monotonicity in the above expressions, implies that we can replace the last term in each of them by r in order to obtain stronger inequalities. Therefore, it suffices for $(r_x)_{x\geq 1}$ and $(\beta_x)_{x\geq 1}$ to satisfy

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

which is equivalent to

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

We seek $(r_x)_{x\geq 1}$ and $(\beta_x)_{x\geq 1}$ in the form $r_x = \gamma^{x-1} + (1-\gamma^{x-1})r$ and $\beta_x = \beta\gamma^{x-1}$ for some $\gamma, \beta \in (0,1)$. Note that this choice ensures Points 1 and 3. Therefore, we just need to show that we can find $\gamma, \beta \in [0,1]$ for which the inequalities in (1) hold for any $s \in S$. Substituting r_x and β_x in terms of γ and β, the inequalities in (1) reduce to

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

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

$$\gamma \geq (1 + \beta rs[\tau + (1-\tau)r])^{-\frac{1}{rs}},$$
$$\gamma \leq (1 - \beta rs[1 - (1-\tau)r])^{\frac{1}{K(1-r)s}}. \tag{2}$$

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

$$\max_{s\in\{s_1,\ldots,s_n\}} (1 + \beta rs[\tau + (1-\tau)r])^{-\frac{1}{rs}} \leq \min_{s\in\{s_1,\ldots,s_n\}} (1 - \beta rs[1 - (1-\tau)r])^{\frac{1}{K(1-r)s}}. \tag{3}$$

Note that the LHS of (3) is monotonically increasing in $s > 0$ whereas the RHS is monotonically decreasing in $s > 0$, therefore it suffices to find $\beta \in (0,1)$ for which

$$(1 + \beta rs_n[\tau + (1-\tau)r])^{-\frac{1}{rs_n}} \leq (1 - \beta rs_1[1 - (1-\tau)r])^{\frac{1}{K(1-r)s_1}}. \tag{4}$$

By Taylor's theorem $(1 + y)^\alpha = 1 + \alpha y + O(y^2)$, so Taylor expanding both sides of (4) in $\beta > 0$ we get

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

Therefore, if we show that $[\tau + (1-\tau)r] > \frac{r}{K(1-r)}[1 - (1-\tau)r]$, the linear coefficient of β on the LHS of (4) will be strictly smaller than the linear coefficient of β on the RHS. Thus, for sufficiently small $\beta > 0$, (4) will hold, which concludes the proof of the lemma. This condition is equivalent to

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

which is true by assumption. Thus, Points 1, 3, and 4 hold. In the full version, we show that Point 2 holds. ◄

4.4   General mean-payoff taxman games

We extend the solution to general games. Recall that the threshold ratio in mean-payoff games is a necessary and sufficient initial ratio with which Max can guarantee a payoff of at least 0.

▶ **Theorem 13.** *Threshold ratios exist in mean-payoff taxman games.*

Proof. Consider a mean-payoff taxman game $\mathcal{G} = \langle V, E, w \rangle$ with taxman parameter τ. If $\mathcal{G}$ is strongly-connected, then by Theorem 9, the threshold ratio in all vertices in $\mathcal{G}$ is the same and is $r \in (0, 1)$ for r such that $\mathtt{MP}(\mathtt{RT}^{F(\tau,r)}(\mathcal{G})) = 0$. If no such r exists, then either $\mathtt{MP}(\mathtt{RT}^{F(\tau,1)}(\mathcal{G})) < 0$, in which case the threshold ratios are 1, or $\mathtt{MP}(\mathtt{RT}^{F(\tau,0)}(\mathcal{G})) > 0$, in which case the threshold ratios are 0. The proof for general games follows along the same lines as the proof for reachability games. For each bottom strongly-connected component S_i of $\mathcal{G}$ we find the threshold ratio $r_i \in (0, 1)$ as in the above. We play a "generalized" reachability game on $\mathcal{G}$ as follows. The game ends once the token reaches one of the BSCCs in $\mathcal{G}$. Max wins the game iff the first time the game enters a BSCC S_i, Max's ratio is greater than r_i. Showing existence of threshold ratios in the generalized game follows the same argument as for reachability games [13].   ◀

5   Computational Complexity

We show, for the first time, computational complexity results for taxman games. We study the following problem, which we call THRESH: given a taxman game $\mathcal{G}$ with taxman parameter τ and a vertex v_0 in $\mathcal{G}$, decide whether $\mathtt{Th}(v_0) \geq 0.5$. The correspondence in Theorem 9 gives the second part of the following theorem, and for the first part, in the full version, we show a reduction from THRESH to the *existential theory of the reals* [7].

▶ **Theorem 14.** *For taxman reachability, parity, and mean-payoff games THRESH is in PSPACE. For strongly-connected mean-payoff games, THRESH is in NP ∩ coNP.*

6   Discussion

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

References

1   M. Aghajohari, G. Avni, and T. A. Henzinger. Determinacy in Discrete-Bidding Infinite-Duration Games. In *In Proc. 30th CONCUR*, 2019.
2   K.R. Apt and E. Grädel. *Lectures in Game Theory for Computer Scientists*. Cambridge University Press, 2011.
3   N. Atzei, M. Bartoletti, and T. Cimoli. A survey of attacks on Ethereum smart contracts. *IACR Cryptology ePrint Archive*, 2016:1007, 2016.

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

7   J. F. Canny. Some Algebraic and Geometric Computations in PSPACE. In *Proc. 20th STOC*, pages 460–467, 1988.

8   K. Chatterjee. Robustness of Structurally Equivalent Concurrent Parity Games. In *Proc. 15th FoSSaCS*, pages 270–285, 2012.

9   K. Chatterjee, A. K. Goharshady, and Y. Velner. Quantitative Analysis of Smart Contracts. In *Proc. 27th ESOP*, pages 739–767, 2018.

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

11 M. Develin and S. Payne. Discrete Bidding Games. *The Electronic Journal of Combinatorics*, 17(1):R85, 2010.

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

13 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.

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

15 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`.

16 S. Muthukrishnan. Ad Exchanges: Research Issues. In *Proc. 5th WINE*, pages 1–12, 2009.

17 N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

18 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.

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

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

21 M.O. Rabin. Decidability of Second Order Theories and Automata on Infinite Trees. *Transaction of the AMS*, 141:1–35, 1969.

22 E. Solan. Continuity of the value of competitive Markov decision processes. *Journal of Theoretical Probability*, 16:831–845, 2003.