

Robustness of Structurally Equivalent Concurrent Parity Games^{*}

Krishnendu Chatterjee

IST Austria (Institute of Science and Technology Austria)

Abstract. We consider two-player stochastic games played on a finite state space for an infinite number of rounds. The games are *concurrent*: in each round, the two players (player 1 and player 2) choose their moves independently and simultaneously; the current state and the two moves determine a probability distribution over the successor states. We also consider the important special case of turn-based stochastic games where players make moves in turns, rather than concurrently. We study concurrent games with ω -regular winning conditions specified as *parity* objectives. The value for player 1 for a parity objective is the maximal probability with which the player can guarantee the satisfaction of the objective against all strategies of the opponent. We study the problem of continuity and robustness of the value function in concurrent and turn-based stochastic parity games with respect to imprecision in the transition probabilities. We present quantitative bounds on the difference of the value function (in terms of the imprecision of the transition probabilities) and show the value continuity for structurally equivalent concurrent games (two games are structurally equivalent if the supports of the transition functions are the same and the probabilities differ). We also show robustness of optimal strategies for structurally equivalent turn-based stochastic parity games. Finally, we show that the value continuity property breaks without the structural equivalence assumption (even for Markov chains) and show that our quantitative bound is asymptotically optimal. Hence our results are tight (the assumption is both necessary and sufficient) and optimal (our quantitative bound is asymptotically optimal).

1 Introduction

Concurrent stochastic games are played by two players on a finite state space for an infinite number of rounds. In every round, the two players simultaneously and independently choose moves (or actions), and the current state and the two chosen moves determine a probability distribution over the successor states. The outcome of the game (or a *play*) is an infinite sequence of states. These games were introduced by Shapley [25], and have been one of the most fundamental and well studied game models in stochastic graph games. We consider ω -regular objectives specified as parity objectives; that is, given an ω -regular set Φ of infinite state sequences, player 1 wins if the outcome of the game lies in Φ . Otherwise, player 2 wins, i.e., the game is zero-sum.

^{*} The research was supported by Austrian Science Fund (FWF) Grant No P 23499-N23, FWF NFN Grant No S11407-N23 (RiSE), ERC Start grant (279307: Graph Games), and Microsoft faculty fellows award.

The class of concurrent stochastic games subsumes many other important classes of games as sub-classes: (1) *turn-based stochastic* games, where in every round only one player chooses moves (i.e., the players make moves in turns); and (2) *Markov decision processes (MDPs)* (one-player stochastic games). Concurrent games and the sub-classes provide a rich framework to model various classes of dynamic reactive systems, and ω -regular objectives provide a robust specification language to express all commonly used properties in verification, and all ω -regular objectives can be expressed as parity objectives. Thus concurrent games with parity objectives provide the mathematical framework to study many important problems in the synthesis and verification of reactive systems [7,24,22] (see also [1,15,2]).

The player-1 *value* $v_1(s)$ of the game at a state s is the limit probability with which player 1 can ensure that the outcome of the game lies in Φ ; that is, the value $v_1(s)$ is the maximal probability with which player 1 can guarantee Φ against all strategies of player 2. Symmetrically, the player-2 *value* $v_2(s)$ is the limit probability with which player 2 can ensure that the outcome of the game lies outside Φ . The problem of studying the computational complexity of MDPs, turn-based stochastic games, and concurrent games with parity objectives has received a lot of attention in literature. Markov decision processes with ω -regular objectives have been studied in [9,10,5] and the results show existence of pure (deterministic) memoryless (stationary) optimal strategies for parity objectives and the problem of value computation is achievable in polynomial time. Turn-based stochastic games with the special case of reachability objectives have been studied in [8] and existence of pure memoryless optimal strategies has been established and the decision problem of whether the value at a state is at least a given rational value lies in $\text{NP} \cap \text{coNP}$. The existence of pure memoryless optimal strategies for turn-based stochastic games with parity objectives was established in [6,29], and again the decision problem lies in $\text{NP} \cap \text{coNP}$. Concurrent parity games have been studied in [11,13,4,16] and for concurrent parity games optimal strategies need not exist, and ε -optimal strategies (for $\varepsilon > 0$) require both infinite memory and randomization in general, and the decision problem can be solved in PSPACE.

Almost all results in the literature consider the problem of computing values and optimal strategies when the game model is given precisely along with the objective. However, it is often unrealistic to know the precise probabilities of transition which are only estimated through observation. Since the transition probabilities are not known precisely, an extremely important question is how robust is the analysis of concurrent games and its sub-classes with parity objectives with respect to small changes in the transition probabilities. This question has been largely ignored in the study of concurrent and turn-based stochastic parity games. In this paper we study the following problems related to continuity and robustness of values: (1) (*continuity of values*): under what conditions can continuity of the value function be proved for concurrent parity games; (2) (*robustness of values*): can quantitative bounds be obtained on the difference of the value function in terms of the difference of the transition probabilities; and (3) (*robustness of optimal strategies*): do optimal strategies of a game remain ε -optimal, for $\varepsilon > 0$, if the transition probabilities are slightly changed.

Our contributions. Our contributions are as follows:

1. We consider *structurally equivalent* game structures, where the supports of the transition probabilities are the same, but the precise transition probabilities may differ. We show the following results for structurally equivalent concurrent parity games:
 - (a) *Quantitative bound.* We present a quantitative bound on the difference of the value functions of two structurally equivalent game structures in terms of the difference of the transition probabilities. We show when the difference in the transition probabilities are small, our bound is asymptotically optimal. Our example to show the matching lower bound is on a Markov chain, and thus our result shows that the bound for a Markov chain can be generalized to concurrent games.
 - (b) *Value continuity.* We show *value continuity* for structurally equivalent concurrent parity games, i.e., as the difference in the transition probabilities goes to 0, the difference in value functions also goes to 0. We then show that the structural equivalence assumption is necessary: we show a family of Markov chains (that are not structurally equivalent) where the difference of the transition probabilities goes to 0, but the difference in the value functions is 1. It follows that the structural equivalence assumption is both necessary (even for Markov chains) and sufficient (even for concurrent games).

It follows from above that our results are both optimal (quantitative bounds) as well as tight (assumption both necessary and sufficient). Our result for concurrent parity games is also a significant quantitative generalization of a result for concurrent parity games of [11] which shows that the set of states with value 1 remains same if the games are structurally equivalent. We also argue that the structural equivalence assumption is not unrealistic in many cases: a reactive system consists of many state variables, and given a state (valuation of variables) it is typically known which variables are possibly updated, and what is unknown is the precise transition probabilities (which are estimated by observation). Thus the system that is obtained for analysis is structurally equivalent to the underlying original system and it only differs in precise transition probabilities.

2. For turn-based stochastic parity games the value continuity and the quantitative bounds are same as for concurrent games. We also prove a stronger result for structurally equivalent turn-based stochastic games that shows that along with continuity of the value function, there is also robustness property for pure memoryless optimal strategies. More precisely, for all $\varepsilon > 0$, we present a bound $\beta > 0$, such that any pure memoryless optimal strategy in a turn-based stochastic parity game is an ε -optimal strategy in every structurally equivalent turn-based stochastic game such that the transition probabilities differ by at most β . Our result has deep significance as it allows the rich literature of work on turn-based stochastic games to carry over robustly for structurally equivalent turn-based stochastic games. As argued before the model of turn-based stochastic game obtained to analyze may differ slightly in precise transition probabilities, and our results shows that the analysis on the slightly imprecise model using the classical results carry over to the underlying original system with small error bounds.

Our results are obtained as follows. The result of [12] shows that the value function for concurrent parity games can be characterized as the limit of the value function of concurrent multi-discounted games (concurrent discounted games with different discount factors associated with every state). There exists bound on difference on value function of discounted games [17], however, the bound depends on the discount factor, and in the limit gives trivial bounds (and in general this approach does not work as value continuity cannot be proven in general and the structural equivalence assumption is necessary). We use a classical result on Markov chains by Friedlin and Wentzell [18] and generalize a result of Solan [26] from Markov chains with single discount to Markov chains with multi-discounted objective to obtain a bound that is independent of the discount factor for structurally equivalent games. Then the bound also applies when we take the limit of the discount factors, and gives us the desired bound.

Our paper is organized as follows: in Section 2 we present the basic definitions, in Section 3 we consider Markov chains with multi-discounted and parity objectives; in Section 4 (Subsection 4.1) we prove the results related to turn-based stochastic games (item (2) of our contributions) and finally in Subsection 4.2 we present the quantitative bound and value continuity for concurrent games along with the two examples to illustrate the asymptotic optimality of the bound and the structural equivalence assumption is necessary. Detailed proofs available in [3].

2 Definitions

In this section we define game structures, strategies, objectives, values and present other preliminary definitions.

Probability Distributions. For a finite set A , a *probability distribution* on A is a function $\delta : A \mapsto [0, 1]$ such that $\sum_{a \in A} \delta(a) = 1$. We denote the set of probability distributions on A by $\mathcal{D}(A)$. Given a distribution $\delta \in \mathcal{D}(A)$, we denote by $\text{Supp}(\delta) = \{x \in A \mid \delta(x) > 0\}$ the *support* of the distribution δ .

Concurrent Game Structures. A (two-player) *concurrent stochastic game structure* $G = \langle S, A, \Gamma_1, \Gamma_2, \delta \rangle$ consists of the following components.

- A finite state space S and a finite set A of moves (or actions).
- Two move assignments $\Gamma_1, \Gamma_2 : S \mapsto 2^A \setminus \emptyset$. For $i \in \{1, 2\}$, assignment Γ_i associates with each state $s \in S$ the nonempty set $\Gamma_i(s) \subseteq A$ of moves available to player i at state s .
- A probabilistic transition function $\delta : S \times A \times A \mapsto \mathcal{D}(S)$, which associates with every state $s \in S$ and moves $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$ a probability distribution $\delta(s, a_1, a_2) \in \mathcal{D}(S)$ for the successor state.

Plays. At every state $s \in S$, player 1 chooses a move $a_1 \in \Gamma_1(s)$, and simultaneously and independently player 2 chooses a move $a_2 \in \Gamma_2(s)$. The game then proceeds to the successor state t with probability $\delta(s, a_1, a_2)(t)$, for all $t \in S$. For all states $s \in S$ and moves $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$, we indicate by $\text{Dest}(s, a_1, a_2) = \text{Supp}(\delta(s, a_1, a_2))$ the set of possible successors of s when moves a_1, a_2 are selected. A *path* or a *play* of G is an infinite sequence $\omega = \langle s_0, s_1, s_2, \dots \rangle$ of states in S

such that for all $k \geq 0$, there are moves $a_1^k \in \Gamma_1(s_k)$ and $a_2^k \in \Gamma_2(s_k)$ such that $s_{k+1} \in \text{Dest}(s_k, a_1^k, a_2^k)$. We denote by Ω the set of all paths. We denote by θ_i the random variable that denotes the i -th state of a path. For a play $\omega = \langle s_0, s_1, s_2, \dots \rangle \in \Omega$, we define $\text{Inf}(\omega) = \{s \in S \mid s_k = s \text{ for infinitely many } k \geq 0\}$ to be the set of states that occur infinitely often in ω .

Special Classes of Games. We consider the following special classes of concurrent games.

1. *Turn-based stochastic games.* A game structure G is *turn-based stochastic* if at every state at most one player can choose among multiple moves; that is, for every state $s \in S$ there exists at most one $i \in \{1, 2\}$ with $|\Gamma_i(s)| > 1$.
2. *Markov decision processes.* A game structure is a *player-1 Markov decision process (MDP)* if for all $s \in S$ we have $|\Gamma_2(s)| = 1$, i.e., only player 1 has choice of actions in the game. Similarly, a game structure is a *player-2 MDP* if for all $s \in S$ we have $|\Gamma_1(s)| = 1$.
3. *Markov chains.* A game structure is a Markov chain if for all $s \in S$ we have $|\Gamma_1(s)| = 1$ and $|\Gamma_2(s)| = 1$. Hence in a Markov chain the players do not matter, and for the rest of the paper a Markov chain consists of a tuple (S, δ) where $\delta : S \mapsto \mathcal{D}(S)$ is the probabilistic transition function.

Strategies. A *strategy* for a player is a recipe that describes how to extend a play. Formally, a strategy for player $i \in \{1, 2\}$ is a mapping $\pi_i : S^+ \mapsto \mathcal{D}(A)$ that associates with every nonempty finite sequence $x \in S^+$ of states, representing the past history of the game, a probability distribution $\pi_i(x)$ used to select the next move. The strategy π_i can prescribe only moves that are available to player i ; that is, for all sequences $x \in S^*$ and states $s \in S$, we require that $\text{Supp}(\pi_i(x \cdot s)) \subseteq \Gamma_i(s)$. We denote by Π_i the set of all strategies for player $i \in \{1, 2\}$.

Given a state $s \in S$ and two strategies $\pi_1 \in \Pi_1$ and $\pi_2 \in \Pi_2$, we define $\text{Outcome}(s, \pi_1, \pi_2) \subseteq \Omega$ to be the set of paths that can be followed by the game, when the game starts from s and the players use the strategies π_1 and π_2 . Formally, $\langle s_0, s_1, s_2, \dots \rangle \in \text{Outcome}(s, \pi_1, \pi_2)$ if $s_0 = s$ and if for all $k \geq 0$ there exist moves $a_1^k \in \Gamma_1(s_k)$ and $a_2^k \in \Gamma_2(s_k)$ such that (i) $\pi_1(s_0, \dots, s_k)(a_1^k) > 0$; (ii) $\pi_2(s_0, \dots, s_k)(a_2^k) > 0$; and (iii) $s_{k+1} \in \text{Dest}(s_k, a_1^k, a_2^k)$. Once the starting state s and the strategies π_1 and π_2 for the two players have been chosen, the probabilities of events are uniquely defined [28], where an *event* $\mathcal{A} \subseteq \Omega$ is a measurable set of paths¹. For an event $\mathcal{A} \subseteq \Omega$, we denote by $\text{Pr}_s^{\pi_1, \pi_2}(\mathcal{A})$ the probability that a path belongs to \mathcal{A} when the game starts from s and the players use the strategies π_1 and π_2 .

Classification of Strategies. We consider the following special classes of strategies.

1. (*Pure*). A strategy π is *pure (deterministic)* if for all $x \in S^+$ there exists $a \in A$ such that $\pi(x)(a) = 1$. Thus, deterministic strategies are equivalent to functions $S^+ \mapsto A$.

¹ To be precise, we should define events as measurable sets of paths *sharing the same initial state*, and we should replace our events with families of events, indexed by their initial state. However, our (slightly) improper definition leads to more concise notation.

2. (*Finite-memory*). Strategies in general are *history-dependent* and can be represented as follows: let M be a set called *memory* to remember the history of plays (the set M can be infinite in general). A strategy with memory can be described as a pair of functions: (a) a *memory update* function $\pi_u : S \times M \mapsto M$, that given the memory M with the information about the history and the current state updates the memory; and (b) a *next move* function $\pi_n : S \times M \mapsto \mathcal{D}(A)$ that given the memory and the current state specifies the next move of the player. A strategy is *finite-memory* if the memory M is finite.
3. (*Memoryless*). A *memoryless* strategy is independent of the history of play and only depends on the current state. Formally, for a memoryless strategy π we have $\pi(x \cdot s) = \pi(s)$ for all $s \in S$ and all $x \in S^*$. Thus memoryless strategies are equivalent to functions $S \mapsto \mathcal{D}(A)$.
4. (*Pure memoryless*). A strategy is *pure memoryless* if it is both pure and memoryless. Pure memoryless strategies neither use memory, nor use randomization and are equivalent to functions $S \mapsto A$.

Qualitative Objectives. We specify *qualitative* objectives for the players by providing the set of *winning plays* $\Phi \subseteq \Omega$ for each player. In this paper we study only zero-sum games [23,17], where the objectives of the two players are complementary. A general class of objectives are the Borel objectives [20]. A *Borel objective* $\Phi \subseteq S^\omega$ is a Borel set in the Cantor topology on S^ω . In this paper we consider ω -regular objectives, which lie in the first $2^{1/2}$ levels of the Borel hierarchy (i.e., in the intersection of Σ_3 and Π_3) [27]. All ω -regular objectives can be specified as parity objectives, and hence in this work we focus on parity objectives, and they are defined as follows.

- *Parity objectives.* For $c, d \in \mathbb{N}$, we let $[c..d] = \{c, c + 1, \dots, d\}$. Let $p : S \mapsto [0..d]$ be a function that assigns a *priority* $p(s)$ to every state $s \in S$, where $d \in \mathbb{N}$. The *Even parity objective* requires that the minimum priority visited infinitely often is even. Formally, the set of winning plays is defined as $\text{Parity}(p) = \{\omega \in \Omega \mid \min(p(\text{Inf}(\omega))) \text{ is even}\}$.

Quantitative Objectives. *Quantitative* objectives are measurable functions $f : \Omega \mapsto \mathbb{R}$. We will consider *multi-discounted* objective functions, as there is a close connection established between concurrent games with multi-discounted objectives and concurrent games with parity objectives. Given a concurrent game structure with state space S , let λ be a *discount vector* that assigns for all $s \in S$ a discount factor $0 < \lambda(s) < 1$ (unless otherwise mentioned we will always consider discount vectors λ such that for all $s \in S$ we have $0 < \lambda(s) < 1$). Let $r : S \mapsto \mathbb{R}$ be a reward function that assigns a real-valued reward $r(s)$ to every state $s \in S$. The multi-discounted objective function $\text{MDT}(\lambda, r) : \Omega \mapsto \mathbb{R}$ maps every path to the mean-discounted reward of the path. Formally, the function is defined as follows: for a path $\omega = s_0 s_1 s_2 \dots$ we have

$$\text{MDT}(\lambda, r)(\omega) = \frac{\sum_{j=0}^{\infty} (\prod_{i=0}^j \lambda(s_i)) \cdot r(s_j)}{\sum_{j=0}^{\infty} (\prod_{i=0}^j \lambda(s_i))}.$$

Also note that a parity objective Φ can be interpreted as a function $\Phi : \Omega \mapsto \{0, 1\}$ by simply considering the characteristic function that assigns 1 to paths that belong to Φ and 0 otherwise.

Values, Optimality, ε -Optimality. Given an objective Φ which is a measurable function $\Phi : \Omega \mapsto \mathbb{R}$, we define the *value* for player 1 of game G with objective Φ from the state $s \in S$ as $\text{Val}(G, \Phi)(s) = \sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2} \mathbb{E}_s^{\pi_1, \pi_2}(\Phi)$; i.e., the value is the maximal expectation with which player 1 can guarantee the satisfaction of Φ against all player 2 strategies. Given a player-1 strategy π_1 , we use the notation $\text{Val}^{\pi_1}(G, \Phi)(s) = \inf_{\pi_2 \in \Pi_2} \mathbb{E}_s^{\pi_1, \pi_2}(\Phi)$. A strategy π_1 for player 1 is *optimal* for an objective Φ if for all states $s \in S$, we have $\text{Val}^{\pi_1}(G, \Phi)(s) = \text{Val}(G, \Phi)(s)$. For $\varepsilon > 0$, a strategy π_1 for player 1 is *ε -optimal* if for all states $s \in S$, we have $\text{Val}^{\pi_1}(G, \Phi)(s) \geq \text{Val}(G, \Phi)(s) - \varepsilon$. The notion of values, optimal and ε -optimal strategies for player 2 are defined analogously. The following theorem summarizes the results in literature related to determinacy and memory complexity of concurrent games and its sub-classes for parity and multi-discounted objectives.

Theorem 1. *The following assertions hold:*

1. (Determinacy [21]). *For all concurrent game structures and for all parity and multi-discounted objectives Φ we have $\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2} \mathbb{E}_s^{\pi_1, \pi_2}(\Phi) = \inf_{\pi_2 \in \Pi_2} \sup_{\pi_1 \in \Pi_1} \mathbb{E}_s^{\pi_1, \pi_2}(\Phi)$.*
2. (Memory complexity). *For all concurrent game structures and for all multi-discounted objectives Φ , randomized memoryless optimal strategies exist [25]. For all turn-based stochastic game structures and for all multi-discounted objectives Φ , pure memoryless optimal strategies exist [17]. For all turn-based stochastic game structures and for all parity objectives Φ , pure memoryless optimal strategies exist [6,29]. In general optimal strategies need not exist in concurrent games with parity objectives, and ε -optimal strategies, for $\varepsilon > 0$, need both randomization and infinite memory in general [11].*

The results of [12] established that the value of concurrent games with certain special multi-discounted objectives can be characterized as valuations of quantitative discounted μ -calculus formula. In the limit, the value function of the discounted μ -calculus formula characterizes the value function of concurrent games with parity objectives. An elegant interpretation of the result was given in [19], and from the interpretation we obtain the following theorem.

Theorem 2 ([12,19]). *Let G be a concurrent game structure with a parity objective Φ defined by a priority function p . Let r be a reward function that assigns reward 1 to even priority states and reward 0 to odd priority states. Then there exists an order $s_1 s_2 \dots s_n$ on the states (where $S = \{s_1, s_2, \dots, s_n\}$) dependent only on the priority function p such that $\text{Val}(G, \Phi) = \lim_{\lambda(s_1) \rightarrow 1} \lim_{\lambda(s_2) \rightarrow 1} \dots \lim_{\lambda(s_n) \rightarrow 1} \text{Val}(G, \text{MDT}(\lambda, r))$; in other words, if we consider the value function $\text{Val}(G, \text{MDT}(\lambda, r))$ with the multi-discounted objective and take the limit of the discount factors to 1 in the order of the states we obtain the value function for the parity objective.*

We now present notions related to *structure equivalent* game structures and distances.

Structure Equivalent Game Structures. Given two game structures $G_1 = \langle S, A, \Gamma_1, \Gamma_2, \delta_1 \rangle$ and $G_2 = \langle S, A, \Gamma_1, \Gamma_2, \delta_2 \rangle$ on the same state and action space, with different transition function, we say that G_1 and G_2 are *structure equivalent*

(denoted $G_1 \equiv G_2$) if for all $s \in S$ and all $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$ we have $\text{Supp}(\delta_1(s, a_1, a_2)) = \text{Supp}(\delta_2(s, a_1, a_2))$. Similarly, two Markov chains $G_1 = (S, \delta_1)$ and $G_2 = (S, \delta_2)$ are structurally equivalent (denoted $G_1 \equiv G_2$) if for all $s \in S$ we have $\text{Supp}(\delta_1(s)) = \text{Supp}(\delta_2(s))$. For a game structure G (resp. Markov chain G), we denote by $\llbracket G \rrbracket_{\equiv}$ the set of all game structures (resp. Markov chains) that are structurally equivalent to G .

Ratio and Absolute Distances. Given two game structures $G_1 = \langle S, A, \Gamma_1, \Gamma_2, \delta_1 \rangle$ and $G_2 = \langle S, A, \Gamma_1, \Gamma_2, \delta_2 \rangle$, the *absolute distance* of the game structures is maximum absolute difference in the transition probabilities. Formally, $\text{dist}_A(G_1, G_2) = \max_{s, t \in S, a \in \Gamma_1(s), b \in \Gamma_2(s)} |\delta_1(s, a, b)(t) - \delta_2(s, a, b)(t)|$. The absolute distance for two Markov chains $G_1 = (S, \delta_1)$ and $G_2 = (S, \delta_2)$ is $\text{dist}_A(G_1, G_2) = \max_{s, t \in S} |\delta_1(s)(t) - \delta_2(s)(t)|$. We now define the ratio distance between two structurally equivalent game structures and Markov chains. Let G_1 and G_2 be two structurally equivalent game structures. The *ratio distance* is defined on the ratio of the transition probabilities. Formally,

$$\text{dist}_R(G_1, G_2) = \max \left\{ \frac{\delta_1(s, a, b)(t)}{\delta_2(s, a, b)(t)}, \frac{\delta_2(s, a, b)(t)}{\delta_1(s, a, b)(t)} \mid s \in S, a \in \Gamma_1(s), b \in \Gamma_2(s), t \in \text{Supp}(\delta_1(s, a, b)) = \text{Supp}(\delta_2(s, a, b)) \right\} - 1$$

The ratio distance between two structurally equivalent Markov chains G_1 and G_2 is $\max \left\{ \frac{\delta_1(s)(t)}{\delta_2(s)(t)}, \frac{\delta_2(s)(t)}{\delta_1(s)(t)} \mid s \in S, t \in \text{Supp}(\delta_1(s)) = \text{Supp}(\delta_2(s)) \right\} - 1$.

Remarks about the Distance Functions. We first remark that the ratio distance is not necessarily a metric. Consider the Markov chain with state space $S = \{s, s'\}$ and let $\varepsilon \in (0, 1/7)$. For $k = 1, 2, 5$ consider the transition functions δ_k such that $\delta_k(t)(s) = 1 - \delta_k(t)(s') = k \cdot \varepsilon$, for all $t \in S$. Let G_k be the Markov chain with transition function δ_k . Then we have $\text{dist}_R(G_1, G_2) = 1$, $\text{dist}_R(G_2, G_5) = \frac{3}{2}$ and $\text{dist}_R(G_1, G_5) = 4$, and hence $\text{dist}_R(G_1, G_2) + \text{dist}_R(G_2, G_5) < \text{dist}_R(G_1, G_5)$. The above example is from [26]. Also note that dist_R is only defined for structurally equivalent game structures, and without the assumption dist_R is ∞ . We also remark that the absolute distance that measures the difference in the transition probabilities is the most intuitive measure for the difference of two game structures.

Proposition 1. *Let G_1 be a game structure (resp. Markov chain) such that the minimum positive transition probability is $\eta > 0$. For all game structures (resp. Markov chains) $G_2 \in \llbracket G_1 \rrbracket_{\equiv}$ we have $\text{dist}_R(G_1, G_2) \leq \frac{\text{dist}_A(G_1, G_2)}{\eta}$.*

Notation for Fixing Strategies. Given a concurrent game structure $G = \langle S, A, \Gamma_1, \Gamma_2, \delta \rangle$, let π_1 be a randomized memoryless strategy. Fixing the strategy π_1 in G we obtain a player-2 MDP, denoted as $G \upharpoonright \pi_1$, defined as follows: (1) the state space is S ; (2) for all $s \in S$ we have $\Gamma_1(s) = \{\perp\}$ (hence it is a player-2 MDP); (3) the new transition function δ_{π_1} is defined as follows: for all $s \in S$ and all $b \in \Gamma_2(s)$ we have $\delta_{\pi_1}(s, \perp, b)(t) = \sum_{a \in \Gamma_1(s)} \pi_1(s)(a) \cdot \delta(s, a, b)(t)$. Similarly if we fix a

randomized memoryless strategy π_1 in an MDP G we obtain a Markov chain, denoted as $G \upharpoonright \pi_1$. The following proposition is straightforward to verify from the definitions.

Proposition 2. *Let G_1 and G_2 be two concurrent game structures (resp. MDPs) that are structurally equivalent. Let π_1 be a randomized memoryless strategy. Then $dist_A(G_1 \upharpoonright \pi_1, G_2 \upharpoonright \pi_1) \leq dist_A(G_1, G_2)$ and $dist_R(G_1 \upharpoonright \pi_1, G_2 \upharpoonright \pi_1) \leq dist_R(G_1, G_2)$.*

3 Markov Chains with Multi-discounted and Parity Objectives

In this section we consider Markov chains with multi-discounted and parity objectives. We present a bound on the difference of value functions of two structurally equivalent Markov chains that is dependent on the distance between the Markov chains and is *independent* of the discount factors. The result for parity objectives is then a consequence of our result for multi-discounted objectives and Theorem 2. Our result crucially depends on a result of Friedlin and Wentzell for Markov chains and we present this result below, and then use it to obtain the main result of the section.

Result of Friedlin and Wentzell. Let (S, δ) be a Markov chain and let s_0 be the initial state. Let $C \subset S$ be a proper subset of S and let us denote by $ex_C = \inf\{n \in \mathbb{N} \mid \theta_n \notin C\}$ the first hitting time to the set $S \setminus C$ of states (or the first exit time from set C) (recall that θ_n is the random variable to denote the n -th state of a path). Let $\mathcal{F}(C, S) = \{f : C \mapsto S\}$ denote the set of all functions from C to S . For every $f \in \mathcal{F}(C, S)$ we define a directed graph $G_f = (S, E_f)$ where $(s, t) \in E_f$ iff $f(s) = t$. Let $\alpha_f = 1$ if the directed graph G_f has no directed cycles (i.e., G_f is a directed acyclic graph); and $\alpha_f = 0$ otherwise. Observe that since f is a function, for every $s \in C$ there is exactly one path that starts at s . For every $s \in C$ and every $t \in S$, let $\beta_f(s, t) = 1$ if the directed path that leaves s in G_f reaches t , otherwise $\beta_f(s, t) = 0$. We now state a result that can be obtained as a special case of the result from Friedlin and Wentzell [18]. Below we use the formulation of the result as presented in [26] (Lemma 2 of [26]).

Theorem 3 (Friedlin-Wentzell result [18]). *Let (S, δ) be a Markov chain, and let $C \subset S$ be a proper subset of S such that $\Pr_s(ex_C < \infty) > 0$ for every $s \in C$ (i.e., from all $s \in C$ with positive probability the first hitting time to the complement set is finite). Then for every initial state $s_1 \in C$ and for every $t \notin C$ we have*

$$\Pr_{s_1}(\theta_{ex_C} = t) = \frac{\sum_{f \in \mathcal{F}(C, S)} (\beta_f(s_1, t) \cdot \prod_{s \in C} \delta(s)(f(s)))}{\sum_{f \in \mathcal{F}(C, S)} (\alpha_f \cdot \prod_{s \in C} \delta(s)(f(s)))}, \tag{1}$$

in other words, the probability that the exit state is t when the starting state is s_1 is given by the expression on the right hand side (very informally the right hand side is the normalized polynomial expression for exit probabilities).

Value Function Difference for Markov Chains. We will use the result of Theorem 3 to obtain bounds on the value functions of Markov chains. We start with the notion of mean-discounted time.

Mean-Discounted Time. Given a Markov chain (S, δ) and a discount vector λ , we define for every state $s \in S$, the *mean-discounted time* the process is in the state s . We first define the mean-discounted time function $\text{MDT}(\lambda, s) : \Omega \mapsto \mathbb{R}$ that maps every path to the mean-discounted time that the state s is visited, and the function is formally defined as follows: for a path $\omega = s_0 s_1 s_2 \dots$ we have

$$\text{MDT}(\lambda, s)(\omega) = \frac{\sum_{j=0}^{\infty} (\prod_{i=0}^j \lambda(s_i)) \cdot \mathbf{1}_{s_j=s}}{\sum_{j=0}^{\infty} (\prod_{i=0}^j \lambda(s_i))};$$

where $\mathbf{1}_{s_j=s}$ is the indicator function. The expected mean-discounted time function for a Markov chain G with transition function δ is defined as follows: $\text{MT}(s_1, s, G, \lambda) = \mathbb{E}_{s_1}[\text{MDT}(\lambda, s)]$, i.e., it is the expected mean-discounted time for s when the starting state is s_1 , where the expectation measure is defined by the Markov chain with transition function δ . We now present a lemma that shows the value function for multi-discounted Markov chains can be expressed as ratio of two polynomials (the result is obtained as a simple extension of a result of Solan [26]).

Lemma 1. *For Markov chains defined on state space S , for all initial states s_0 , for all states s , for all discount vectors λ , there exists two polynomials $g_1(\cdot)$ and $g_2(\cdot)$ in $|S|^2$ variables $x_{t,t'}$, where $t, t' \in S$ such that the following conditions hold:*

1. *the polynomials have degree at most $|S|$ with non-negative coefficients; and*
2. *for all transition functions δ over S we have $\text{MT}(s_0, s, G, \lambda) = \frac{g_1(\delta)}{g_2(\delta)}$, where $G = (S, \delta)$, $g_1(\delta)$ and $g_2(\delta)$ denote the values of the function g_1 and g_2 such that all the variables $x_{t,t'}$ is instantiated with values $\delta(t)(t')$ as given by the transition function δ .*

Proof. (Sketch). We present a sketch of the proof (details in [3]). Fix a discount vector λ . We construct a Markov chain $\bar{G} = (\bar{S}, \bar{\delta})$ as follows: $\bar{S} = S \cup S_1$, where S_1 is a copy of states of S (and for a state $s \in S$ we denote its corresponding copy as s_1); and the transition function $\bar{\delta}$ is defined below

1. $\bar{\delta}(s_1)(s_1) = 1$ for all $s_1 \in S_1$ (i.e., all copy states are absorbing);
2. for $s \in S$ we have

$$\bar{\delta}(s)(t) = \begin{cases} (1 - \lambda(s)) & t = s_1; \\ \lambda(s) \cdot \delta(s)(t) & t \in S; \\ 0 & t \in S_1 \setminus s_1; \end{cases}$$

i.e., it goes to the copy with probability $(1 - \lambda(s))$, it follows the transition δ in the original copy with probabilities multiplied by $\lambda(s)$.

We first show that for all s_0 and s we have $\text{MT}(s_0, s, G, \lambda) = \text{Pr}_{s_0}^{\bar{\delta}}(\theta_{\text{ex}_S} = s_1)$; i.e., the expected mean-discounted time in s when the original Markov chain starts in s_0 is the probability in the Markov chain $(\bar{S}, \bar{\delta})$ that the first hitting state out of S is the copy s_1 of the state s . The claim is easy to verify as both $(\text{MT}(s_0, s, G, \lambda))_{s_0 \in S}$

and $(\text{Pr}_{s_0}^{\bar{\delta}}(\theta_{\text{ex}_S} = s_1))_{s_0 \in S}$ are the unique solution of the following system of linear equations: for all $t \in S$ we have $y_t = (1 - \lambda(t)) \cdot \mathbf{1}_{t=s} + \sum_{z \in S} \lambda(t) \cdot \delta(t)(z) \cdot y_z$.

We now claim that $\text{Pr}_{s_0}^{\bar{\delta}}(\text{ex}_S < \infty) > 0$ for all $s_0 \in S$. This follows since for all $s \in S$ we have $\bar{\delta}(s)(s_1) = (1 - \lambda(s)) > 0$ and since $s_1 \notin S$ we have $\text{Pr}_{s_0}^{\bar{\delta}}(\text{ex}_S = 2) = (1 - \lambda(s_0)) > 0$. Now we observe that we can apply Theorem 3 on the Markov chain $\bar{G} = (\bar{S}, \bar{\delta})$ with S as the set C of states of Theorem 3, and obtain the result. Indeed the terms α_f and $\beta_f(s, t)$ are independent of δ , and the two products of Equation (1) each contains at most $|S|$ terms of the form $\bar{\delta}(s)(t)$ for $s, t \in \bar{S}$. Thus the desired result follows. ■

Lemma 2. *Let $h(x_1, x_2, \dots, x_k)$ be a polynomial function with non-negative coefficients of degree at most n . Let $\varepsilon > 0$ and $y, y' \in \mathbb{R}^k$ be two non-negative vectors such that for all $i = 1, 2, \dots, k$ we have $\frac{1}{1+\varepsilon} \leq \frac{y_i}{y'_i} \leq 1 + \varepsilon$. Then we have $(1 + \varepsilon)^{-n} \leq \frac{h(y)}{h(y')} \leq (1 + \varepsilon)^n$.*

Lemma 3. *Let $G_1 = (S, \delta)$ and $G_2 = (S, \delta')$ be two structurally equivalent Markov chains. For all non-negative reward functions $r : S \mapsto \mathbb{R}$ such that the reward function is bounded by 1, for all discount vectors λ , for all $s \in S$ we have $|\text{Val}(G_1, \text{MDT}(\lambda, r))(s) - \text{Val}(G_2, \text{MDT}(\lambda, r))(s)| \leq (1 + \text{dist}_R(G_1, G_2))^{2 \cdot |S|} - 1$; i.e., the absolute difference of the value functions for the multi-discounted objective is bounded by $(1 + \text{dist}_R(G_1, G_2))^{2 \cdot |S|} - 1$.*

The proof of Lemma 3 uses Lemma 1 and Lemma 2 and details available in [3].

Theorem 4. *Let $G_1 = (S, \delta)$ and $G_2 = (S, \delta')$ be two structurally equivalent Markov chains. Let η be the minimum positive transition probability in G_1 . The following assertions hold:*

1. *For all non-negative reward functions $r : S \mapsto \mathbb{R}$ such that the reward function is bounded by 1, for all discount vectors λ , for all $s \in S$ we have*

$$|\text{Val}(G_1, \text{MDT}(\lambda, r))(s) - \text{Val}(G_2, \text{MDT}(\lambda, r))(s)| \leq (1 + \varepsilon_R)^{2 \cdot |S|} - 1$$

$$\leq (1 + \varepsilon_A)^{2 \cdot |S|} - 1$$

2. *For all parity objectives Φ and for all $s \in S$ we have*

$$|\text{Val}(G_1, \Phi)(s) - \text{Val}(G_2, \Phi)(s)| \leq (1 + \varepsilon_R)^{2 \cdot |S|} - 1 \leq (1 + \varepsilon_A)^{2 \cdot |S|} - 1$$

where $\varepsilon_R = \text{dist}_R(G_1, G_2)$ and $\varepsilon_A = \frac{\text{dist}_A(G_1, G_2)}{\eta}$.

Proof. The first part follows from Lemma 3 and Proposition 1. The second part follows from part 1, the fact the value function for parity objectives is obtained as the limit of multi-discounted objectives (Theorem 2), and the fact the bound for part 1 is independent of the discount factors (hence independent of taking the limit). ■

Remark on Structural Assumption in the Proof. The result of the previous theorem depends on the structural equivalence assumption in two crucial ways. They are as

follows: (1) Proposition 1 that establishes the relation of $dist_R$ and $dist_A$ only holds with the assumption of structural equivalence; and (2) without the structural equivalence assumption $dist_R$ is ∞ , and hence without the assumption the bound of the previous theorem is ∞ , which is a trivial bound. We will later show (in Example 1) that the structural equivalence assumption is necessary.

4 Value Continuity for Parity Objectives

In this section we show two results: first we show robustness of strategies and present quantitative bounds on value functions for turn-based stochastic games and then we show continuity for concurrent parity games.

4.1 Bounds for Structurally Equivalent Turn-Based Stochastic Parity Games

In this section we present quantitative bounds for robustness of optimal strategies in structurally equivalent turn-based stochastic games. For every $\varepsilon > 0$, we present a bound $\beta > 0$, such that if the distance of the structurally equivalent turn-based stochastic games differs by at most β , then any pure memoryless optimal strategy in one game is ε -optimal in the other. The result is first shown for MDPs and then extended to turn-based stochastic games (both proofs are in [3]).

Theorem 5. *Let G_1 be a turn-based stochastic game such that the minimum positive transition probability is $\eta > 0$. The following assertions hold:*

1. *For all turn-based stochastic games $G_2 \in \llbracket G_1 \rrbracket_{\equiv}$, for all parity objectives Φ and for all $s \in S$ we have*

$$\begin{aligned} |\text{Val}(G_1, \Phi)(s) - \text{Val}(G_2, \Phi)(s)| &\leq (1 + dist_R(G_1, G_2))^{2 \cdot |S|} - 1 \\ &\leq \left(1 + \frac{dist_A(G_1, G_2)}{\eta}\right)^{2 \cdot |S|} - 1 \end{aligned}$$

2. *For $\varepsilon > 0$, let $\beta \leq \frac{\eta}{2} \cdot \left(1 + \frac{\varepsilon}{2}\right)^{\frac{1}{2 \cdot |S|}} - 1$. For all $G_2 \in \llbracket G_1 \rrbracket_{\equiv}$ such that $dist_A(G_1, G_2) \leq \beta$, for all parity objectives Φ , every pure memoryless optimal strategy π_1 in G_1 is an ε -optimal strategy in G_2 .*

4.2 Value Continuity for Concurrent Parity Games

In this section we show value continuity for structurally equivalent concurrent parity games, and show with an example on Markov chains that the continuity property breaks without the structural equivalence assumption. Finally with an example on Markov chains we show the our quantitative bounds are asymptotically optimal for small distance values. We start with a lemma for MDPs.

Lemma 4. *Let G_1 and G_2 be two structurally equivalent MDPs. Let η be the minimum positive transition probability in G_1 . For all non-negative reward functions $r : S \mapsto \mathbb{R}$ such that the reward function is bounded by 1, for all discount vectors λ , for all $s \in S$ we have*

$$\begin{aligned} |\text{Val}(G_1, \text{MDT}(\lambda, r))(s) - \text{Val}(G_2, \text{MDT}(\lambda, r))(s)| &\leq (1 + \text{dist}_R(G_1, G_2))^{2 \cdot |S|} - 1 \\ &\leq \left(1 + \frac{\text{dist}_A(G_1, G_2)}{\eta}\right)^{2 \cdot |S|} - 1 \end{aligned}$$

The main idea of the proof of the above lemma is to fix a pure memoryless optimal strategy and then use the results for Markov chains. Using the same proof idea, along with randomized memoryless optimal strategies for concurrent game structures and the above lemma, we obtain the following lemma (the result is identical to the previous lemma, but for concurrent game structures instead of MDPs).

Lemma 5. *Let G_1 and G_2 be two structurally equivalent concurrent game structures. Let η be the minimum positive transition probability in G_1 . For all non-negative reward functions $r : S \mapsto \mathbb{R}$ such that the reward function is bounded by 1, for all discount vectors λ , for all $s \in S$ we have*

$$\begin{aligned} |\text{Val}(G_1, \text{MDT}(\lambda, r))(s) - \text{Val}(G_2, \text{MDT}(\lambda, r))(s)| &\leq (1 + \text{dist}_R(G_1, G_2))^{2 \cdot |S|} - 1 \\ &\leq \left(1 + \frac{\text{dist}_A(G_1, G_2)}{\eta}\right)^{2 \cdot |S|} - 1 \end{aligned}$$

We now present the main theorem that depends on Lemma 5.

Theorem 6. *Let G_1 and G_2 be two structurally equivalent concurrent game structures. Let η be the minimum positive transition probability in G_1 . For all parity objectives Φ and for all $s \in S$ we have*

$$\begin{aligned} |\text{Val}(G_1, \Phi)(s) - \text{Val}(G_2, \Phi)(s)| &\leq (1 + \text{dist}_R(G_1, G_2))^{2 \cdot |S|} - 1 \\ &\leq \left(1 + \frac{\text{dist}_A(G_1, G_2)}{\eta}\right)^{2 \cdot |S|} - 1 \end{aligned}$$

Proof. The result follows from Theorem 2, Lemma 5 and the fact that the bound of Lemma 5 are independent of the discount factors and hence independent of taking the limits. ■

In the following theorem we show that for structurally equivalent game structures, for all parity objectives, the value function is continuous in the absolute distance between the game structures. We have already remarked (after Theorem 4) that the structural equivalence assumption is required in our proofs, and we show in Example 1 that this assumption is necessary.

Theorem 7. *For all concurrent game structures G_1 , for all parity objectives Φ*

$$\lim_{\varepsilon \rightarrow 0} \sup_{G_2 \in \llbracket G_1 \rrbracket_{\equiv, \text{dist}_A(G_1, G_2) \leq \varepsilon}} \sup_{s \in S} |\text{Val}(G_1, \Phi)(s) - \text{Val}(G_2, \Phi)(s)| = 0.$$

Proof. Let $\eta > 0$ be the minimum positive transition probability in G_1 . By Theorem 6 we have

$$\lim_{\varepsilon \rightarrow 0} \sup_{G_2 \in \llbracket G_1 \rrbracket_{\equiv}, \text{dist}_A(G_1, G_2) \leq \varepsilon} \sup_{s \in S} |\text{Val}(G_1, \Phi)(s) - \text{Val}(G_2, \Phi)(s)| \leq \lim_{\varepsilon \rightarrow 0} \left(1 + \frac{\varepsilon}{\eta}\right)^{2 \cdot |S|} - 1$$

The above limit equals to 0, and the desired result follows. \blacksquare

Example 1 (Structurally equivalence assumption necessary). In this example we show that in Theorem 7 the structural equivalence assumption is necessary, and thereby show that the result is tight. We show an Markov chain G_1 and a family of Markov chains G_2^ε , for $\varepsilon > 0$, such that $\text{dist}_A(G_1, G_2^\varepsilon) \leq \varepsilon$ (but G_1 is not structurally equivalent to G_2^ε) with a parity objective Φ and we have $\lim_{\varepsilon \rightarrow 0} \sup_{s \in S} |\text{Val}(G_1, \Phi)(s) - \text{Val}(G_2^\varepsilon, \Phi)(s)| = 1$. The Markov chains G_1 and G_2^ε are defined over the state space $\{s_0, s_1\}$, and in G_1 both states have self-loops with probability 1, and in G_2^ε the self-loop at s_0 has probability $1 - \varepsilon$ and the transition probability from s_0 to s_1 is ε (for details see [3]). Clearly, $\text{dist}_A(G_1, G_2^\varepsilon) = \varepsilon$. The parity objective Φ requires to visit the state s_1 infinitely often (i.e., assign priority 2 to s_1 and priority 1 to s_0). Then we have $\text{Val}(G_1, \Phi)(s_0) = 0$ as the state s_0 is never left, whereas in G_2^ε the state s_1 is the only closed recurrent set of the Markov chain and hence reached with probability 1 from s_0 . Hence $\text{Val}(G_2^\varepsilon, \Phi)(s_0) = 1$. It follows that $\lim_{\varepsilon \rightarrow 0} \sup_{s \in S} |\text{Val}(G_1, \Phi)(s) - \text{Val}(G_2^\varepsilon, \Phi)(s)| = 1$. \blacksquare

Example 2 (Asymptotically tight bound for small distances). We now show that our quantitative bound for the value function difference is asymptotically optimal for small distances. Let us denote the absolute distance as ε , and the quantitative bound we obtain in Theorem 6 is $(1 + \frac{\varepsilon}{\eta})^{2 \cdot |S|} - 1$, and if ε is small, then we obtain the following approxi-

mate bound: $\left(1 + \frac{\varepsilon}{\eta}\right)^{2 \cdot |S|} - 1 \approx 1 + 2 \cdot |S| \cdot \frac{\varepsilon}{\eta} - 1 = 2 \cdot |S| \cdot \frac{\varepsilon}{\eta}$. We now illustrate with an example (on structurally equivalent Markov chains) where the difference in the value function is $O(|S| \cdot \varepsilon)$, for small ε . Consider the Markov chain defined on state space $S = \{s_0, s_1, \dots, s_{2n-1}, s_{2n}\}$ as follows: states s_0 and s_{2n} are absorbing (states with self-loops of probability 1) and for a state $1 \leq i \leq 2n-1$ we have $\delta(s_i)(s_{i-1}) = \frac{1}{2} + \varepsilon$; and $\delta(s_i)(s_{i+1}) = \frac{1}{2} - \varepsilon$; i.e., we have a Markov chain defined on a line from 0 to $2n$ (with 0 and $2n$ absorbing states) and the chain moves towards 0 with probability $\frac{1}{2} + \varepsilon$ and towards $2n$ with probability $\frac{1}{2} - \varepsilon$ (for complete details see [3]). Our goal is to estimate the probability to reach the state s_0 , and let v_i denote the probability to reach s_0 from the starting state s_i . We show (details in [3]) that if $\varepsilon = 0$, then $v_n = \frac{1}{2}$ and for $0 < \varepsilon < \frac{1}{2}$, such that ε is close to 0, we have $v_n = \frac{1}{2} + n \cdot \varepsilon$. Observe that the Markov chain obtained for $\varepsilon = 0$ and $\frac{1}{2} > \varepsilon > 0$ are structurally equivalent. Thus the desired result follows. \blacksquare

5 Conclusion

In this work we studied the robustness and continuity property of concurrent and turn-based stochastic parity games with respect to small imprecision in the transition probabilities. We presented (i) quantitative bounds on difference of the value functions and

proved value continuity for concurrent parity games under the structural equivalence assumption, and (ii) showed robustness of all pure memoryless optimal strategies for structurally equivalent turn-based stochastic parity games. We also showed that the structural equivalence assumption is necessary and that our quantitative bounds are asymptotically optimal for small imprecision. We believe our results will find applications in robustness analysis of various other classes of stochastic games.

References

1. Abadi, M., Lamport, L., Wolper, P.: Realizable and Unrealizable Specifications of Reactive Systems. In: Ronchi Della Rocca, S., Ausiello, G., Dezanı-Ciancaglini, M. (eds.) ICALP 1989. LNCS, vol. 372, pp. 1–17. Springer, Heidelberg (1989)
2. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. *Journal of the ACM* 49, 672–713 (2002)
3. Chatterjee, K.: Robustness of structurally equivalent concurrent parity games. *CoRR* 1107.2009 (2011)
4. Chatterjee, K., de Alfaro, L., Henzinger, T.A.: The complexity of quantitative concurrent parity games. In: SODA 2006, pp. 678–687. ACM-SIAM (2004)
5. Chatterjee, K., Henzinger, T.A., Jurdziński, M.: Games with secure equilibria. In: LICS 2004, pp. 160–169. IEEE (2004)
6. Chatterjee, K., Jurdziński, M., Henzinger, T.A.: Quantitative stochastic parity games. In: SODA 2004, pp. 121–130. SIAM (2004)
7. Church, A.: Logic, arithmetic, and automata. In: *Proceedings of the International Congress of Mathematicians*, pp. 23–35. Institut Mittag-Leffler (1962)
8. Condon, A.: The complexity of stochastic games. *Information and Computation* 96(2), 203–224 (1992)
9. Courcoubetis, C., Yannakakis, M.: The complexity of probabilistic verification. *Journal of the ACM* 42(4), 857–907 (1995)
10. de Alfaro, L.: *Formal Verification of Probabilistic Systems*. PhD thesis. Stanford University (1997)
11. de Alfaro, L., Henzinger, T.A.: Concurrent omega-regular games. In: LICS 2000, pp. 141–154. IEEE (2000)
12. de Alfaro, L., Henzinger, T.A., Majumdar, R.: Discounting the Future in Systems Theory. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 1022–1037. Springer, Heidelberg (2003)
13. de Alfaro, L., Majumdar, R.: Quantitative solution of omega-regular games. In: STOC 2001, pp. 675–683. ACM Press (2001)
14. Derman, C.: *Finite State Markovian Decision Processes*. Academic Press (1970)
15. Dill, D.L.: *Trace Theory for Automatic Hierarchical Verification of Speed-independent Circuits*. The MIT Press (1989)
16. Etesami, K., Yannakakis, M.: Recursive Concurrent Stochastic Games. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006, Part II. LNCS, vol. 4052, pp. 324–335. Springer, Heidelberg (2006)
17. Filar, J., Vrieze, K.: *Competitive Markov Decision Processes*. Springer, Heidelberg (1997)
18. Friedlin, M.I., Wentzell, A.D.: *Random perturbations of dynamical systems*. Springer, Heidelberg (1984)
19. Gimbert, H., Zielonka, W.: Discounting infinite games but how and why? *Electr. Notes Theor. Comput. Sci.* 119(1), 3–9 (2005)
20. Kechris, A.: *Classical Descriptive Set Theory*. Springer, Heidelberg (1995)

21. Martin, D.A.: The determinacy of Blackwell games. *The Journal of Symbolic Logic* 63(4), 1565–1581 (1998)
22. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: *POPL 1989*, pp. 179–190. ACM Press (1989)
23. Raghavan, T.E.S., Filar, J.A.: Algorithms for stochastic games — a survey. *ZOR — Methods and Models of Op. Res.* 35, 437–472 (1991)
24. Ramadge, P.J., Wonham, W.M.: Supervisory control of a class of discrete-event processes. *SIAM Journal of Control and Optimization* 25(1), 206–230 (1987)
25. Shapley, L.S.: Stochastic games. *Proc. Nat. Acad. Sci. USA* 39, 1095–1100 (1953)
26. Solan, E.: Continuity of the value of competitive Markov decision processes. *Journal of Theoretical Probability* 16, 831–845 (2003)
27. Thomas, W.: Automata on infinite objects. In: van Leeuwen, J. (ed.) *Handbook of Theoretical Computer Science*, vol. B, pp. 133–191. Elsevier (1990)
28. Vardi, M.Y.: Automatic verification of probabilistic concurrent finite-state systems. In: *FOCS 1985*, pp. 327–338. IEEE Computer Society Press (1985)
29. Zielonka, W.: Perfect-Information Stochastic Parity Games. In: Walukiewicz, I. (ed.) *FOSSACS 2004*. LNCS, vol. 2987, pp. 499–513. Springer, Heidelberg (2004)