

Obligation Blackwell Games and p-Automata

Krishnendu Chatterjee
Institute of Science and Technology Austria

Nir Piterman
University of Leicester

November 5, 2013

Abstract

We recently introduced p-automata, automata that read discrete-time Markov chains and showed they provide an automata-theoretic framework for reasoning about pCTL model checking and abstraction of discrete time Markov chains. We used turn-based stochastic parity games to define acceptance of Markov chains by a special subclass of p-automata. Definition of acceptance required a reduction to a series of turn-based stochastic parity games. The reduction was cumbersome and complicated and could not support acceptance by general p-automata, which was left undefined as there was no notion of games that supported it.

Here we generalize two-player games by adding a structural acceptance condition called *obligations*. Obligations are orthogonal to the linear winning conditions that define whether a play is winning. Obligations are a declaration that player 0 can achieve a certain value from a configuration. If the obligation is met, the value of that configuration for player 0 is 1.

One cannot define value in obligation games by the standard mechanism of considering the measure of winning paths on a Markov chain and taking the supremum of the infimum of all strategies. Mainly because obligations need definition even for Markov chains and the nature of obligations has the flavor of an infinite nesting of supremum and infimum operators. We define value via a reduction to turn-based games similar to Martin's proof of determinacy of Blackwell games with Borel objectives. Based on this value definition we show that obligation games are determined. We show that for Markov chains with Borel objectives and obligations, and finite turn-based stochastic parity games with obligations there exists an alternative and simpler characterization of the value function without going through a Martin-like reduction. Based on this simpler definition we give an exponential time algorithm to analyze finite turn-based stochastic parity games with obligations and show that the decision problem of winning parity games with obligations is in $NP \cap co-NP$. Finally, we show that obligation games provide the necessary framework for reasoning about p-automata and that they generalize the previous definition.

1 Introduction

Markov chains are a very important modeling formalism in many areas of science. In computer science, Markov chains form the basis of central techniques such as performance modeling, and the design and correctness of randomized algorithms used in security and communication protocols. Recognizing this prominent role of Markov chains, the formal-methods community has devoted significant attention to these models, e.g., in developing model checking for *qualitative* [11, 6, 27] and *quantitative* [1] properties, logics for reasoning about Markov chains [10, 17], and probabilistic simulation and bisimulation [18, 17]. Model-checking tools such as PRISM [13] and LiQuor [5] support such reasoning about Markov chains and have users in many fields of computer science and beyond.

The automata-theoretic approach to verification has proven to be very powerful for reasoning about systems modeled as Kripke structures. For example, it supports algorithms for satisfiability of temporal logics [7], model checking [16], and abstraction [12].

We recently introduced p-automata, which are devices that read Markov chains as input [14]. We showed that p-automata provide an automata-theoretic framework for reasoning about pCTL model checking, and abstraction of discrete time Markov chains. The definition of p-automata is motivated by pCTL [9], the de-facto standard for model checking Markov chains, and alternating tree automata: they combine the rich combinatorial structure of alternating automata with pCTL's ability to quantify the probabilities of regular sets of paths. Acceptance of Kripke structures by alternating tree automata is decided by solving turn-based games (cf. [8]). Similarly, acceptance of Markov chains by p-automata is decided by solving turn-based *stochastic* games. However, acceptance of p-automata was defined through a complicated and cumbersome reduction to solving a *series* of turn-based stochastic parity games. Furthermore, this reduction supported only a subclass of p-automata, which we called *uniform*, and could not be generalized to unrestricted p-automata. Intuitively, uniform automata separate measuring probability of regular path sets and setting thresholds on these probabilities. For uniform p-automata, we showed how acceptance can be decided by a series of turn-based stochastic parity games. Acceptance for general p-automata could not be defined as there was no game framework that supported the unbounded interaction between measuring probability and setting thresholds.

Here, we propose a game notion that supports such interaction. In order to do that we augment winning conditions in games by adding a structural acceptance condition called *obligations*. A winning condition is a combination of a classical set of winning paths and obligations on some of the game configurations. An obligation is a declaration by Player 0 that she can win with a certain value from the configuration. Then, in order to be able to derive a non-zero value from a configuration with obligation, Player 0 has to ensure that the measure of paths that satisfy the winning condition from that configuration meets her promised value. If she can do that, the configuration will have the value 1 for her. That is, the value of the configuration does not depend on the measure of winning paths obtained in the game following the visit to that configuration, only on whether the obligation is met or not. In other words, in order to meet an obligation the measure of the union of the following sets of paths must satisfy the




Figure 1: Intuition regarding obligations. Circles denote states of Markov chains. The name of the state is written in the top half and the obligation (if exists) is written in the bottom half. Probabilities are written next to edges except in case there is a single outgoing edge whose probability is 1.




Figure 2: A set of measure 0 matters.

value constraint of the obligation: paths that (i) reach other obligations that can be met, or (ii) paths that never reach other obligations and satisfy the winning condition. In addition, paths that visit infinitely many obligations must satisfy the winning condition as well.

Consider the example in Figure 1. Suppose that every path that visits s_2 infinitely often is rejecting and every accepting path that visits s_4 infinitely often must visit s_3 infinitely often. In addition, s_2 has an obligation of more than $\frac{2}{3}$. That is, in order for s_2 to have a positive value, the measure of the set of paths starting in s_2 that satisfies the acceptance condition must be more than $\frac{2}{3}$. If the obligation is met, the value of s_2 is 1. Otherwise, the value of s_2 is 0. Similarly, s_3 has an obligation of at least $\frac{1}{2}$. The obligation at s_2 cannot be met. Every set of paths starting at s_2 with measure more than $\frac{2}{3}$ must contain a path that reaches s_2 again. Thus, in order to fulfill the obligation of s_2 , it has to be visited again, recursively, where the recursion is unfounded. As the path that visits s_2 infinitely often is rejecting, s_2 's obligation cannot be met. Thus, s_2 's value is 0. On the other hand, the obligation of s_3 can be met. Indeed, the measure of paths that start in s_3 and reach s_3 again without passing through s_2 (i.e., $\neg s_2$ until s_3) is exactly $\frac{1}{2}$. Whenever a path reaches s_3 the same obligation needs to be met again. So the same set of paths is used again. Let us consider the set of infinite paths obtained by concatenating infinitely many finite segments (from obligations to obligations, i.e., concatenating paths from s_3 to itself without visiting s_2) considered above. These paths visit s_3 infinitely often and never visit s_2 , and thus satisfy the winning condition.

When obligations are involved, the value depends on sets of measure 0. Consider, for example, the Markov chain in Figure 2. Suppose that the set of winning paths includes all paths. Then, the obligation of configuration s_2 can be met. Indeed, the probability to reach another obligation that can be met (s_2) or never reach obligations and win is 1. It follows that the value of configuration s_1 is 1. However, by removing from the set of winning paths the single path $(s_1 \cdot s_2)^\omega$, whose measure is 0, this

changes. The obligation at configuration s_2 can no longer be met. Every set of paths starting from s_2 of measure more than half must contain a path that reaches s_2 again. Thus, in order to meet the obligation of s_2 , the path $(s_1 \cdot s_2)^\omega$ must be included. This is a losing path that makes infinitely many obligations. Hence, the value of s_2 is 0 and the value of s_1 is $\frac{1}{2}$.

These two examples consider the most simple types of interaction, where no player makes choices and everything is determined by chance. Blackwell games with Borel objectives are a very general form of graph games. Blackwell games are two-player games where both players choose simultaneously and independently their actions (also known as concurrent games). From every configuration a choice of actions determines a distribution over successor configurations. Borel objectives are Borel sets of infinite paths of configurations. We now discuss the main challenges involved in giving a formal definition of the value in general obligation Blackwell games. The typical way of defining a value in a two-player game is done by a sequence of steps: First, fixing strategies for the two players we get a Markov chain for which the measure of winning paths is well defined. Then, fixing a strategy for Player 0 one can get a value of this strategy by considering the infimum over all strategies for Player 1. Finally, the value of Player 0 in the game is the supremum of the values of all her strategies. All three steps fail for obligations. First, it is not even clear how to define the value of a Markov chain, and the value depends on sets of measure 0 (see Figure 2). Then, it is not clear that a strategy for Player 0 has a well defined value. Such a well defined value would depend on independence between the choice of “met obligations” and the choice of strategy for Player 1. Finally, the nature of obligations means that taking the supremum over strategies of Player 0 does not lead to the real value of the game. This follows because in some types of Borel games (without obligations) the game value cannot be attained. There is no single strategy that can achieve the value of the game but an infinite sequence of strategies can get arbitrarily close to the value. A single obligation game requires a player to commit to achieving values possibly infinitely many times. Each obligation could be the result of an infinite sequence of strategies. Hence, a single external supremum does not capture this and we need an infinite nesting of supremum (and infimum) operators.

In order to define the value of Blackwell games with Borel objectives and obligations we use a reduction to turn-based games with Borel objectives similar to Martin’s proof that Blackwell games with Borel objectives are determined [20]. Intuitively, we explicitly add the value to the game and make Player 0 prove that a value can be won by showing how it propagates under probabilistic choices made by both players. We add a new crucial component that captures obligations to Martin’s proof. This ingredient, which we call *concession*, captures the notion that the value for Player 0 may not be achievable but may be approximated arbitrarily close. To capture approximation Player 1 chooses a small concession to grant Player 0 after which Player 0 should not have a problem to show that she can win the required amount. This reduction defines a value for Blackwell Borel obligation games.

A major issue is then that of well definedness of obligation games. Well definedness, or determinacy, means that whatever Player 0 cannot avoid losing Player 1 ensures to win and vice versa. Formally, it says that the sum of values in a game for Player 0 and Player 1 is always 1. This is a fundamental property that needs to be established for

games. For almost all types of games used in verification determinacy has never been an issue. This relies on Martin’s foundational result that Blackwell games with Borel objectives are determined [20]. Blackwell games are general enough so that a simple reduction to them suffices to show determinacy for almost all types of two-player games. However, obligations take our games out of scope of Martin’s result. This is apparent as inclusion/removal of a 0-measure set can change the value of a game. In order to show determinacy of Blackwell games with Borel objectives and obligations we analyze our reduction to turn-based games with Borel objectives. We show that indeed our games are determined.

Our reduction gives the general definition of values and its analysis gives us the determinacy result. However, the reduction is not amenable to computational analysis as it constructs an uncountable game. For computational analysis, we consider Markov chains with Borel objectives and obligations and finite turn-based stochastic parity games with obligations. We show that in these cases, we can embed the notion of winning into the structure of the game by using *choice sets*. Intuitively, these are the obligations that have value 1, i.e., where the obligation of Player 0 is actually met. This gives rise to a simpler definition where we contrast a strategy of one player with the strategy of the other player as customary in definition of games. We show that this simpler definition, which does not work for the general case, coincides with the definition arising from the Martin-like reduction for Markov chains and finite turn-based stochastic parity games with obligations. Based on this direct characterization, we give algorithms that analyze finite turn-based stochastic parity games with obligations. We show how to decide whether the value in such a game is at least (or more) than a given value $r \in [0, 1]$ in $\text{NP} \cap \text{co-NP}$ and to compute the value in exponential time. The algorithm identifies a general choice set and calls a solver for finite turn-based stochastic parity games to check the sanity of the choice set. Our $\text{NP} \cap \text{co-NP}$ bound matches the bounds for the special cases of turn-based stochastic reachability games (without obligations).

We also show that if games with obligations have a finite number of exchanges between obligations and no-obligations, then the analysis of the game can be reduced to the analysis of a series of Blackwell games (with no obligations).

Finally, we return to p-automata and using turn-based stochastic parity games with obligations we define acceptance of general p-automata. We show that the new definition using the obligation games generalizes acceptance of uniform p-automata as defined in [14].

Related works. While we considered an automata theoretic approach to capture pCTL an alternative approach is to consider probabilistic μ -calculus. The problem of considering a probabilistic μ -calculus framework to capture pCTL was considered in [22, 21]. They add an independent choice to turn-based stochastic games and show their determinacy and that such games give a semantics to a probabilistic μ -calculus. In [23], they show decidability of a fragment of their probabilistic μ -calculus in 3EXPTIME. In contrast, our determinacy result for obligation games is for the general class of Blackwell (concurrent) games with Borel objectives, and our decidability result for turn-based stochastic obligation parity games establishes a $\text{NP} \cap \text{co-NP}$ bound. Comparison of the two types of games does not seem simple. Our framework for turn-based stochastic obligation parity games could also provide better algorithmic analysis for fragments of

probabilistic μ -calculus of [22, 21].

2 Background

For a countable set S let $\mathcal{D}(S) = \{d : S \rightarrow [0, 1] \mid \exists T \subseteq S \text{ such that } |T| \in \mathbb{N}, \forall s \notin T. d(s) = 0 \text{ and } \sum_{s \in T} d(s) = 1\}$ be the set of discrete probability distributions with finite support over S . A distribution d is *pure* if there is some $s \in S$ such that $d(s) = 1$.

A *countable labeled Markov chain* M over set of atomic propositions \mathbb{AP} is a tuple (S, P, L, s^{in}) , where S is a countable set of *locations*, $P : S \rightarrow \mathcal{D}(S)$ is a probabilistic transition, $s^{\text{in}} \in S$ the *initial* location, and $L : S \rightarrow 2^{\mathbb{AP}}$ a *labeling function* with $L(s)$ the set of propositions true in location s . We sometimes also treat P as a function $P : S \times S \rightarrow [0, 1]$, where $P(s, s')$ is $P(s)(s')$. Let $\text{succ}(s)$ be the set $\{s' \in S \mid P(s, s') > 0\}$ of *successors* of s . By definition all Markov chains we consider are *finitely branching*, i.e. $\text{succ}(s)$ is finite for all $s \in S$. We write $\text{MC}_{\mathbb{AP}}$ for the set of all (finitely branching) Markov chains over \mathbb{AP} . A *path* π from location s in M is an infinite sequence of locations $s_0 s_1 \dots$ with $s_0 = s$ and $P(s_i, s_{i+1}) > 0$ for all $i \geq 0$.

Given a Markov chain M with set of states S , an open set in S^ω is a set $\{w\} \cdot S^\omega$ for some $w \in S^*$. A set is *Borel* if it is in the σ -algebra defined by these open sets. The measure of every Borel set α is defined as usual in this σ -algebra [2, 25]. We denote the measure of a Borel set α as $\text{Prob}_M(\alpha)$.

Blackwell Games A *Blackwell game* is $G = (V, A_0, A_1, R, \alpha)$, where V is a countable set of configurations, A_0 and A_1 are finite sets of actions, α is a Borel set defining the winning set of Player 0, and $R : V \times A_0 \times A_1 \rightarrow \mathcal{D}(V)$ is a transition function associating with a configuration v and a pair of actions for both players a distribution over next configurations with *finite support*. A play is an infinite sequence $p = v_0 v_1 \dots$ such that for every $i \geq 0$ there are $a_i^0 \in A_0$ and $a_i^1 \in A_1$ such that $R(v_i, a_i^0, a_i^1)(v_{i+1}) > 0$.

A *strategy* for Player 0 is $\sigma : V^+ \rightarrow \mathcal{D}(A_0)$. A strategy for Player 1 is similar. A strategy is *memoryless* if for every $w, w' \in V^*$ and $v \in V$ we have $\sigma(wv) = \sigma(w'v)$ and it is *pure* if for every $w \in V^+$ we have $\sigma(w)$ is pure. Let Σ (resp. Π) be the set of all strategies for Player 0 (resp. Player 1).

Each $(\sigma, \pi) \in \Sigma \times \Pi$ from game G and configuration v determine a Markov chain with locations V^+ . Formally, $v(\sigma, \pi) = (V^+, P, L, v)$, where the labeling function L is irrelevant, and for every $w \in V^*$ and $v' \in V$ we set $P(vwv') = \sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sigma(vwv')(a_0) \cdot \pi(vwv')(a_1) \cdot R(v', a_0, a_1)$.

Sometimes, we may want to start a game from an initial sequence of configurations, which we call *play prefix* or just *prefix*. Let $w = v_0 \dots v_n \in V^+$ be a prefix. Then $w(\sigma, \pi)$ is the Markov chain $(\{w\} \cdot V^*, P, L, w)$, where $P(wuv) = \sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sigma(wuv)(a_0) \cdot \pi(wuv)(a_1) \cdot R(v, a_0, a_1)$, for $u \in V^*$ and $v \in V$. All definitions, generalize to this setting.

The value of (σ, π) for Player 0 from prefix $w \in \{v\} \cdot V^*$, is $\text{Prob}_{w(\sigma, \pi)}(\{w\} \cdot V^\omega \cap \alpha)$, denoted $\text{val}_0(v(\sigma, \pi), w)$. The value of w for Player 0 in G is $\sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \text{val}_0(w(\sigma, \pi), w)$,

denoted $\text{val}_0(G, w)$. Dually, the value of w for Player 1 in G , denoted $\text{val}_1(G, w)$, is $\sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} (1 - \text{val}_0(w(\sigma, \pi), w))$.

Theorem 1 *Let G be a game and α a Borel set. Then for every $w \in V^+$ we have $\text{val}_0(G, w) + \text{val}_1(G, w) = 1$ [20].*

The value $\text{val}_1(G, w)$ can be also obtained by considering the game $\text{dual}(G) = (V, A_1, A_0, \text{dual}(R), V^\omega \setminus \alpha)$, where $\text{dual}(R)(v', a_1, a_0) = R(v', a_0, a_1)$. Formally, $\text{val}_1(G, w) = \text{val}_0(\text{dual}(G), w)$. By definition, for every Markov chain M and every measurable set $\alpha \subseteq V^\omega$ we have $\text{Prob}_M(\alpha) = 1 - \text{Prob}_M(V^\omega \setminus \alpha)$.

Turn-Based Stochastic Games A turn-based stochastic game G is $((V, E), (V_0, V_1, V_p), \kappa, \alpha)$ with the following components.

- V is a countable set of configurations.
- $E \subseteq V^2$ is a set of edges such that for every $v \in V$ we have $|\{v' \mid (v, v') \in E\}|$ is finite.
- The triplet (V_0, V_1, V_p) partitions V so that V_0 is the set of Player 0 configurations, V_1 is the set of Player 1 configurations, and V_p is the set of probabilistic configurations.
- $\kappa : V_p \rightarrow \mathcal{D}(V)$ is such that $\kappa(v)(v') > 0$ if and only if $(v, v') \in E$.
- α is a Borel set as before.

A play is an infinite sequence $v_0 v_1 \dots$ such that for all $i \in \mathbb{N}$ we have $(v_i, v_{i+1}) \in E$. A strategy for Player 0 is a function $\sigma : V^* \cdot V_0 \rightarrow \mathcal{D}(V)$ such that for all $w \in V^*$ and $v \in V_0$ we have $\sigma(wv)(v') > 0$ implies $(v, v') \in E$. Strategies for Player 1 are defined analogously. The type of strategy is determined by the type of game and no confusion will arise. As before $(\sigma, \pi) \in \Sigma \times \Pi$ determine a Markov chain $w(\sigma, \pi)$. Then, the value of Player 0 from prefix w is $\text{val}_0(G, w) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \text{Prob}_{w(\sigma, \pi)}(\alpha)$ and the value of Player 1 from prefix w is $\text{val}_1(G, w) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} (1 - \text{Prob}_{w(\sigma, \pi)}(\alpha))$.

A turn-based stochastic game can be seen as a Blackwell game whose configurations are of the following types:

- v is a Player 0 configuration if for every $a_0 \in A_0$ and $a_1, a'_1 \in A_1$ we have $R(v, a_0, a_1) = R(v, a_0, a'_1)$ and $R(v, a_0, a_1)$ is pure.
- v is a Player 1 configuration if for every $a_0, a'_0 \in A_0$ and $a_1 \in A_1$ we have $R(v, a_0, a_1) = R(v, a'_0, a_1)$ and $R(v, a_0, a_1)$ is pure.
- v is a probabilistic configuration if for every $a_0, a'_0 \in A_0$ and $a_1, a'_1 \in A_1$ we have $R(v, a_0, a_1) = R(v, a'_0, a'_1)$.

Corollary 1 *Let G be a turn-based stochastic game, α a Borel set, and $w \in V^+$. Then, $\text{val}_0(G, w) = 1 - \text{val}_1(G, w)$.*

When $V_p = \emptyset$ the game is *simple* or just *turn based*. For turn-based games the set V does not have to be countable. In this case it is enough to consider pure strategies, which implies that a pair $(\sigma, \pi) \in \Sigma \times \Pi$ induces a unique play $w(\sigma, \pi)$. Then, the value of Player 0 from prefix w is either 1 or 0. Equivalently, Player 0 wins from w if there is a strategy σ such that for every strategy π we have $w(\sigma, \pi) \in \alpha$. Otherwise, Player 1 wins from w . In this case, we write $W_0 = \{w \mid \text{Player 0 wins from } w\}$ and $W_1 = \{w \mid \text{Player 1 wins from } w\}$.

Theorem 2 *Let G be a turn-based game and α a Borel set. Then $W_0 \cap W_1 = \emptyset$ and $W_0 \cup W_1 = V^+$. [19]*

When $w \in W_0$ we also write $\text{val}_0(G, w) = 1$ and $\text{val}_1(G, w) = 0$. Dually, when $w \in W_1$ we write $\text{val}_0(G, w) = 0$ and $\text{val}_1(G, w) = 1$.

We say that α is derived from a parity condition $c : V \rightarrow [0..k]$ if for every play $p = v_0v_1v_2\cdots \in V^\omega$ we have $p \in \alpha$ iff $\liminf_{n \rightarrow \infty} c(v_n)$ is even.

Theorem 3 *Consider a finite game G , where α is derived from a parity condition, $\bowtie \in \{>, \geq\}$, and r is a rational.*

- *If G is a Blackwell game, whether $\text{val}_0(G, w) \bowtie r$ and $\text{val}_1(G, w) \bowtie r$ can be decided in PSPACE [3].*
- *If G is a turn-based stochastic game, the values $\text{val}_0(G, w)$ and $\text{val}_1(G, w)$ can be computed in exponential time and whether $\text{val}_i(G, w) \bowtie r$ can be decided in $\text{NP} \cap \text{co-NP}$ [4].*
- *If G is a turn-based stochastic game, there is a memoryless strategy σ achieving G 's value [4]. That is:*

$$\inf_{\pi \in \Pi} \text{Prob}_{w(\sigma, \pi)}(\alpha) = \text{val}_0(G, w)$$

- *If G is a turn-based game, whether $w \in W_0$ can be decided in $\text{UP} \cap \text{co-UP}$ [15].*
- *If G is a countable turn-based game then there are pure-memoryless strategies σ and π such that σ is winning from every configuration $v \in W_0$ and π is winning from every configuration $v \in W_1$ [26].*

3 Obligation Blackwell Games

We introduce obligation Blackwell games. These games extend Blackwell games by having a winning condition that includes a winning set (as in normal Blackwell games) and a set of obligations. Intuitively, a play is winning for Player 0 if it belongs to the winning set. However, whenever meeting an obligation, Player 0 has to make sure that the value of the game in that configuration satisfies the obligation. If the obligation can be met, the value for Player 0 at the configuration is 1.

An obligation Blackwell game (OBG for short) is $G = (V, A_0, A_1, R, \mathcal{G})$, where V , A_0 , A_1 , and R are like in Blackwell games. The goal $\mathcal{G} = \langle \alpha, O \rangle$, where $\alpha \subseteq V^\omega$ is a Borel set as for Blackwell games and $O : V \rightarrow (\{\geq, >\} \times [0, 1]) \cup \{\perp\}$. The obligation function O associates with some configurations the value \perp saying that there is no obligation associated with this configuration. With other configurations O associates an obligation $>r$ or $\geq r$ stating that Player 0 can use this configuration (i.e., she derives a non-zero value when getting to this configuration and this non-zero value is 1) only if she can ensure that the value she can get from this configuration onwards meets the obligation. It follows that, recursively, Player 0 has to ensure that every obligation configuration satisfies the obligation requirement with plays in α . For configuration v , if $O(v) \neq \perp$ we call v an *obligation* configuration and if $O(v) = \perp$ we call v a *non-obligation* configuration.

As mentioned, the usual approach to defining values in games by considering the measure of winning paths on a Markov chain and taking the supremum of infimum of

$\hat{V} =$	$(V^+ \times (0, 1]) \cup (V^+ \times (0, 1] \times \{\epsilon\}) \cup \left\{ (w, r, f) \mid \sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}(w)} d_0(a_0) \cdot d_1(a_1) \cdot R(v, a_0, a_1)(v') \cdot f(v') > r \right\}$
$V_0 =$	$\{(w, r) \mid \text{either } O(w) = \perp \text{ or } \exists r'. O(w) = >r'\} \cup \{(w, r, \epsilon)\}$
$V_1 =$	$\{(w, r) \mid \exists r'. O(w) = \geq r'\} \cup \{(w, r, f)\}$
$E =$	$\{((w, r), (w, r'', \epsilon)) \mid O(w) = \geq r' \text{ and } 0 < r'' < r'\} \cup \{((w, r, f), (w \cdot v', f(v')) \mid f(v') > 0\} \cup \{((w, r), (w, r', f)) \mid \text{either } O(w) = \perp \text{ and } r' = r \text{ or } O(w) = >r'\} \cup \{((w, r, \epsilon), (w, r, f))\}$
$\hat{\alpha} =$	$\{p \in \hat{V}^\omega \mid p \downarrow_V \in \alpha\}$, where $p \downarrow_V$ is the limit of the projection of p on V^+ .

Figure 3: Components of $\text{turn}(G)$.

strategies of the respective players does not work. This is mainly for two reasons. First, the definition of value over a Markov chain needs defining in its own right. Second, if a value is not achievable by a single strategy (which is the case in Blackwell games with Borel objectives *without* obligations) the supremum over the value of strategies is not sufficient to capture the complexity of obligations and (infinitely many) nested supremum (and infimum) operators are required. Intuitively, the value of a configuration in an obligation game is the value in the modified game where Player 0's objective is to either reach obligations she can fulfil or never reach obligations and fulfil the Borel winning conditions. If during this interaction a new obligation is met then this obligation needs to be fulfilled in the same way. If infinitely many obligations are met along a path, this path has to be winning according to the Borel objective. We present the formal definition of the value for the players in an OBG through a reduction to a turn-based game similar to Martin's proof that Blackwell games are determined [20].

We generalize the function O to apply to prefixes, where $O(wv) = O(v)$ for every $wv \in V^+$ and similarly for R , and succ . Consider the game $\text{turn}(G) = ((\hat{V}, E), (V_0, V_1), \hat{\alpha})$, where the components of $\text{turn}(G)$ are given in Figure 3.

There are three types of configurations. Configurations of the form (w, r) , where $O(w) = \perp$, are illustrated on the left in Figure 4. Such configurations are Player 0 configurations, where she claims that the value of prefix w is *more than* r . From such configurations Player 0 chooses a successor configuration (w, r, f) , where f is a function associating a value to every successor of w that proves that indeed the value at w is greater than r . Configurations of the form (w, r) , where $O(w) = >r'$, are illustrated in the middle in Figure 4. Such configurations are Player 0 configurations, where, ignoring the value r , she has to prove that the value is greater than r' . Thus, she proceeds as above but for the value r' instead of r . Configurations of the form (w, r) , where $O(w) = \geq r'$, are illustrated on the right in Figure 4. Such configurations are Player 1 configurations, where, acknowledging that it may be impossible for Player 0 to achieve exactly r' but possible to achieve every $r'' < r'$, Player 1 grants Player 0 a concession and moves to a configuration (w, r'', ϵ) from which, as above, Player 0 chooses a successor configuration (w, r'', f) . Notice, that ϵ is used as a syntactic symbol signifying that a concession has been granted, it is not a value. Then from configurations of the form (w, r, f) , Player 1 chooses which successor v' of w to follow and proceeds to $(w \cdot v', f(v'))$. Finally, we note that as α is a Borel set, then $\hat{\alpha}$ is also a Borel set.




Figure 4: The structure of $\text{turn}(G)$. Diamonds are Player 0 configurations and rectangles are Player 1 configurations. Shaded areas represent a continuum of edges, where every edge is associated with an entry from the continuous domain written next to the edge. Fans of discrete edges represent finite choice, where every edge is associated with a value from the domain written next to the edge. A dashed octagon is either a Player 0 or Player 1 configuration depending on $O(v')$.

By Theorem 2 for every prefix w and for every value $r \in (0, 1]$ from configuration (w, r) in the game $\text{turn}(G)$ either Player 0 wins or else Player 1 wins.

Lemma 1 *For every OBG G and every prefix w , if Player 0 wins from (w, r) in $\text{turn}(G)$, she wins from every configuration (w, r') for $r' < r$. If Player 1 wins from (w, r) in $\text{turn}(G)$, she wins from every configuration (w, r') for $r' > r$.*

Proof: This can be done by reusing the strategy in $\text{turn}(G)$. Essentially, in order to show $r' < r$ Player 0 can show r . Dually, in order to show that $r' > r$ is infeasible it is enough to show that r is infeasible. \square

So winning values for Player 0 are downward closed and winning values for Player 1 are upward closed. It follows that there is a unique value below which Player 0 wins and above which Player 1 wins.

Corollary 2 *For every OBG G and every prefix w , there is a value $s(G, w) \in [0, 1]$ such that for every $r' < s(G, w)$ Player 0 wins from (w, r') in $\text{turn}(G)$ and for every $r'' > s(G, w)$ Player 1 wins from (w, r'') in $\text{turn}(G)$.*

Notice that Player 0 may or may not win from $s(G, w)$. For a prefix w we define the value of w in G as follows. If $O(w) = \perp$ then the value of w in G , denoted $\text{val}_0(G, w)$, is $s(G, w)$. If $O(w) = \bowtie r$ then $\text{val}_0(G, w)$ is 1 iff $s(G, w) > r$ or $s(G, w) = r$ and $\bowtie = \geq$ and it is 0 otherwise.

We now turn to the issue of determinacy. In order to show that the value of Player 0 and Player 1 sum to 1, we define the dual game. Dualization of a game consists of changing the roles of the two players and switching the goal to the complement. Here, the complementation of the goal is slightly more complicated than usual. Consider a game $G = (V, A_0, A_1, R, \mathcal{G})$, where $\mathcal{G} = \langle \varphi, O \rangle$. The dual game $\text{dual}(G) = (V, A_1, A_0, \text{dual}(R), \text{dual}(\mathcal{G}))$, where $\text{dual}(R)(v, a_1, a_0) = R(v, a_0, a_1)$, $\text{dual}(\mathcal{G}) = \langle V^\omega \setminus \varphi, \text{dual}(O) \rangle$, and $\text{dual}(O)$ is defined below.

$$\text{dual}(O)(v) = \begin{cases} \perp & \text{If } O(v) = \perp \\ >1-r & \text{If } O(v) = \geq r \\ \geq 1-r & \text{If } O(v) = >r \end{cases}$$

Intuitively, if in G Player 0 has the obligation to achieve more than r with the set φ , then the dual player (Player 0 in $\text{dual}(G)$) has the obligation to achieve at least $1-r$ with the goal set $V^\omega \setminus \varphi$. Syntactically, $\text{dual}(\text{dual}(G)) = G$. We use the dual game to define the value for Player 1. Formally, let $\text{val}_1(G, w)$ denote the value of w in $\text{dual}(G)$. We prove that obligation games are determined by showing that the sum of values of a prefix w in G and in $\text{dual}(G)$ is 1.

Lemma 2 *For every OBG and every prefix w , there are distributions $d_0 \in \mathcal{D}(A_0)$ and $d_1 \in \mathcal{D}(A_1)$ such that*

$$s(G, w) = \sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}(v)} d_0(a_0) \cdot d_1(a_1) \cdot R(w, a_0, a_1)(v') \cdot \text{val}_0(G, w \cdot v').$$

Furthermore, for every $d'_0 \in \mathcal{D}(A_0)$ and $d'_1 \in \mathcal{D}(A_1)$ the following hold.

$$\begin{aligned} s(G, w) &\leq \sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}(v)} d_0(a_0) \cdot d'_1(a_1) \cdot R(w, a_0, a_1)(v') \cdot \text{val}_0(G, w \cdot v'). \\ s(G, w) &\geq \sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}(v)} d'_0(a_0) \cdot d_1(a_1) \cdot R(w, a_0, a_1)(v') \cdot \text{val}_0(G, w \cdot v'). \end{aligned}$$

Proof: Consider the values $\text{val}_0(G, w \cdot v')$ for $v' \in \text{succ}(w)$. By Von Neumann's minimax theorem [24] there is an $r \in [0, 1]$ and optimal distributions $d_0 \in \mathcal{D}(A_0)$ and $d_1 \in \mathcal{D}(A_1)$ such that

$$\sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}w} d_0(a_0) \cdot d_1(a_1) \cdot R(w, a_0, a_1)(v') \text{val}_0(G, w \cdot v') = r$$

and for every $d'_0 \in \mathcal{D}(A_0)$ and every $d'_1 \in \mathcal{D}(A_1)$ we have

$$\begin{aligned} \sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}w} d_0(a_0) \cdot d'_1(a_1) \cdot R(w, a_0, a_1)(v') \text{val}_0(G, w \cdot v') &\geq r \\ \sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}w} d'_0(a_0) \cdot d_1(a_1) \cdot R(w, a_0, a_1)(v') \text{val}_0(G, w \cdot v') &\leq r \end{aligned}$$

We have to show that $r = s(G, w)$.

Suppose that $r > s(G, w)$. Let $\delta = r - s(G, w)$ and consider a play starting from $(w, s(G, w) + \frac{\delta}{2}, f)$, where f is the function that associates $\max(0, \text{val}_0(G, w \cdot v') - \frac{\delta}{4})$ to every successor v' of w . Clearly, the minimax value for f is at least $r - \frac{\delta}{4}$, which is larger than $s(G, w) + \frac{\delta}{2}$. By definition of $\text{val}_0(G, w \cdot v')$, Player 0 has a winning strategy from $(w \cdot v', \text{val}_0(G, w \cdot v') - \frac{\delta}{4})$. It follows that Player 0 also wins from $(w, s(G, w) + \frac{\delta}{2})$ contradicting the definition of $s(G, w)$.

Suppose that $r < s(G, w)$. Let $\delta = s(G, w) - r$ and consider a play starting from $(w, s(G, w) - \frac{\delta}{2})$. Consider the configuration $(w, s(G, w) - \frac{\delta}{2})$. In order to win, Player 0 has to choose a successor configuration $(w, s(G, w) - \frac{\delta}{2}, f)$, where f associates at least $\text{val}_0(G, w \cdot v') + \frac{\delta}{2}$ with some successor $w \cdot v'$ of w . Then, by definition of $s(G, w)$, Player 1 wins from configuration $(w \cdot v', f(v'))$ contradicting the definition of $s(G, w)$. \square

We are now ready to prove that value is well defined and that obligation Blackwell games are determined.

Theorem 4 *For All prefixes w in an OBG G we have $\text{val}_0(G, w) + \text{val}_1(G, w) = 1$.*

The proof of Theorem 4 is non-trivial. The proof requires Martin's determinacy proof style analysis of the uncountable game $\text{turn}(G)$, along with new subtleties (for example as shown in the example in Figure 2 that measure zero sets could play an important role in values of obligation games).

Proof: Note that the definition of our values is through turn-based deterministic games, and thus relies on determinacy of turn-based deterministic games. In the present proof we do not explicitly rely on Borel objectives, but the definition of values through turn-based deterministic games requires determinacy for them (and determinacy holds for turn-based deterministic games with Borel objectives). More explicitly, our proof relies on determinacy for turn-based deterministic games rather than Borel objectives. The determinacy proof of Martin also relies on determinacy of turn-based deterministic games.

We add a few comments for readers familiar with Martin's work. We note that Martin considers a quantitative objectives that map plays to payoffs in the range $[0, 1]$ while we consider whether Player 0 is winning or not. This is equivalent to restricting the payoffs to the range $\{0, 1\}$. Furthermore, he uses the symbol for integration to represent the value while we use the notation $\text{val}(\cdot, \cdot)$ and talk about winning. The first part of the proof below corresponds to the construction of the strategy for Player 0 (p. 1570) and the proof of Lemma 1.1 in Martin's paper. The second part of the proof below corresponds to the construction of the strategy for Player 1 (p. 1572) and the proof of Lemma 1.4. The second half of Martin's paper considers various extensions of his result. We do not touch upon similar subjects to his.

For a prefix w , let $S(G, w)$ denote the set of values r such that Player 0 wins from (w, r) in $\text{turn}(G)$.

\Rightarrow We show that if $r \in S(G, w)$ then $1 - r \notin S(\text{dual}(G), w)$.

Suppose that $r \in S(G, w)$. That is, Player 0 wins from (w, r) in game $\text{turn}(G)$.

We show that Player 1 wins from $(w, 1 - r)$ in $\text{turn}(\text{dual}(G))$ proving that $1 -$

$r \notin S(\text{dual}(G), w)$. Let σ be the winning strategy of Player 0 in $\text{turn}(G)$. We now construct a winning strategy for Player 1 in $\text{turn}(\text{dual}(G))$. To distinguish between a prefix of a play in G and prefixes in $\text{turn}(G)$ or $\text{turn}(\text{dual}(G))$ we call the latter two *paths*. For a path $\text{dual}(p)$ in $\text{turn}(\text{dual}(G))$ we use the strategy σ to construct a path p in $\text{turn}(G)$ such that whenever $\text{dual}(p)$ ends in configuration (w', t) then p ends in configuration (w', r) such that $r + t \geq 1$. Initially, we start from configuration $(w, 1 - r)$ in $\text{turn}(\text{dual}(G))$ and from configuration (w, r) in $\text{turn}(G)$. That is, both paths are of length one. Suppose that the paths p and $\text{dual}(p)$ end in configurations (w', r') and (w', t') , respectively, and that $t' + r' \geq 1$. We have the following cases.

- Suppose that $O(w') = \perp$ then (w', r') is a Player 0 configuration in $\text{turn}(G)$ and (w', t') is a Player 0 configuration in $\text{turn}(\text{dual}(G))$. The winning strategy σ instructs Player 0 to choose some configuration (w', r', f) in $\text{turn}(G)$. Suppose that Player 0 chooses the configuration (w', t', f') in $\text{turn}(\text{dual}(G))$. By definition, there has to be a configuration $v' \in \text{succ}(w')$ such that $f(v') + f'(v') \geq 1$. We make Player 1 choose $(w' \cdot v', f(v'))$ in $\text{turn}(G)$ and extend the strategy π of Player 1 in $\text{turn}(\text{dual}(G))$ by choosing $(w' \cdot v', f'(v'))$.
- Suppose that $O(w') = >r''$ in G . Then $O(w') = \geq 1 - r''$ in $\text{dual}(G)$. It follows that (w', r') is a Player 0 configuration in $\text{turn}(G)$ and (w', t') is a Player 1 configuration in $\text{turn}(\text{dual}(G))$. The winning strategy σ instructs us to choose a configuration (w', r'', f) in $\text{turn}(G)$. From the minimax theorem [24] it follows that there is a value $r''' > r''$ that is attained for the optimal choice $d_0 \in \mathcal{D}(A_0)$ such that

$$\inf_{d_1 \in \mathcal{D}(A_1)} \left(\sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}(w')} d_0(a_0) \cdot d_1(a_1) \cdot R(w', a_0, a_1)(v') \cdot f(v') \right)$$

is at least r''' . Let $\delta = r''' - r''$. Notice that $1 - r'' - \delta = 1 - r'''$. Then, from configuration (w', t') in $\text{turn}(\text{dual}(G))$, Player 1 chooses the successor configuration $(w', 1 - r''', \epsilon)$, in effect giving up δ for Player 0's benefit. Suppose that Player 0 chooses the successor configuration $(w', 1 - r''', f')$ in $\text{turn}(\text{dual}(G))$. As above, there has to be a successor $v' \in \text{succ}(w')$ such that $f'(v') + f(v') \geq 1$. Then we make Player 1 choose $(w' \cdot v', f(v'))$ in $\text{turn}(G)$ and extend Player 1's strategy in $\text{turn}(\text{dual}(G))$ by the choice $(w \cdot v', f'(v'))$.

- Suppose that $O(w') = \geq r''$ in G . Then $O(w') = > 1 - r''$ in $\text{dual}(G)$. It follows that (w', r') is a Player 1 configuration in $\text{turn}(G)$ and (w', t') is a Player 0 configuration in $\text{turn}(\text{dual}(G))$. Suppose that Player 0 chooses the successor configuration $(w', 1 - r'', f')$ in $\text{turn}(\text{dual}(G))$. From the minimax theorem [24] it follows that there is a value $r''' < r''$ that is attained for the optimal choice $d_1 \in \mathcal{D}(A_1)$ such that

$$\inf_{d_0 \in \mathcal{D}(A_0)} \left(\sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}(w')} d_0(a_0) \cdot d_1(a_1) \cdot R(w', a_0, a_1)(v') \cdot f'(v') \right)$$

is at least $1 - r''$. Let $\delta = r'' - r'''$. Notice that $r'' - \delta = r'''$. Then, from configuration (w', r') in $\text{turn}(G)$, we make Player 1 choose the successor configuration (w', r''', ϵ) , in effect giving up δ for Player 0's benefit. Now, Player 0's winning strategy in $\text{turn}(G)$ instructs her to choose a configuration (w', r''', f) . As above, there has to be a successor $v' \in \text{succ}(w')$ such that $f(v') + f'(v') \geq 1$. Then we make Player 1 choose $(w' \cdot v', f(v'))$ in $\text{turn}(G)$ and extend Player 1's strategy in $\text{turn}(\text{dual}(G))$ by the choice $(w' \cdot v', f'(v'))$.

Consider the two infinite plays played in $\text{turn}(G)$ and $\text{turn}(\text{dual}(G))$. Clearly, when projecting the two plays on the configurations in $V^+ \times (0, 1]$ that appear in them and then on the configurations in V^+ we get exactly the same play. By assumption σ is a winning strategy for Player 0 in $\text{turn}(G)$. Hence, the limit of this projection is in α implying that the strategy constructed for Player 1 in $\text{turn}(\text{dual}(G))$ is indeed winning.

\Leftarrow We show that $1 - s(G, w) \leq s(\text{dual}(G), w)$. Notice that if $s(G, w) = 1$ then clearly, $1 - s(G, w) \leq s(\text{dual}(G), w)$. We consider the case that $s(G, w) < 1$. By Lemma 1 for every $r > s(G, w)$ we have Player 1 wins in $\text{turn}(G)$ from (w, r) . If $t = 1 - r$ then $t < 1 - s(G, w)$. We show that Player 0 wins from (w, t) in $\text{turn}(\text{dual}(G))$.

Consider some value $r > s(G, w)$ such that Player 1 wins from (w, r) . We show that Player 0 wins from $(w, 1 - r)$ in $\text{turn}(\text{dual}(G))$ by proving that $1 - r \in S(\text{dual}(G), w)$. We use the difference between $1 - r + s(G, w)$ and 1 to give a winning strategy for Player 0 in $\text{turn}(\text{dual}(G))$. We use a winning strategy π of Player 1 in $\text{turn}(G)$ to produce a winning strategy for Player 0 in $\text{turn}(\text{dual}(G))$. For a path $\text{dual}(p)$ in $\text{turn}(\text{dual}(G))$ we use the winning strategy π of Player 1 in $\text{turn}(G)$ to construct a path p in $\text{turn}(G)$ such that whenever $\text{dual}(p)$ ends in configuration (w', t') then p ends in configuration (w', r') such that $r' \geq s(G, w')$, Player 1 is winning from p using π , and $t' < 1 - r'$.

Consider a configuration (w, r) such that $r > s(G, w)$. As $r > s(G, w)$ there is some $r > \tilde{r} > s(G, w)$ such that Player 1 wins from (w, \tilde{r}) . Let π be the winning strategy of Player 1 from (w, \tilde{r}) . Initially, we start from configuration $(w, 1 - r)$ in $\text{turn}(\text{dual}(G))$ and from configuration (w, \tilde{r}) in $\text{turn}(G)$. Clearly, (w, \tilde{r}) is winning for Player 1, $\tilde{r} \geq s(G, w)$, and $1 - r < 1 - \tilde{r}$.

Suppose that the two paths p and $\text{dual}(p)$ end in a configurations (w', r') and (w', t') , respectively, and that $r' \geq s(G, w')$, Player 1 wins from p using π , and $t' < 1 - r'$. We have the following cases.

- Suppose that $O(w') = \perp$ then (w', r') is a Player 0 configuration in $\text{turn}(G)$ and (w', t') is a Player 0 configuration in $\text{turn}(\text{dual}(G))$.

For every location $v' \in \text{succ}(w')$ let $u(v')$ be the following value:

$$\inf\{1, f(v') \mid (w' \cdot v'', r', f) \in \hat{V} \text{ and } \pi(p \cdot (w' \cdot v'', r', f)) = v'\}$$

That is, we consider all possible choices for Player 0 from (w', r') . Such a choice includes a function $f : \text{succ}(w') \rightarrow [0, 1]$. Then, whenever the winning strategy of Player 1 chooses to proceed to v' , we record the value promised by Player 0 and take the infimum of all these values.

By the minimax theorem there are $d_0 \in \mathcal{D}(A_0)$ and $d_1 \in \mathcal{D}(A_1)$ such that

$$\sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v' \in \text{succ}(w')} d_0(a_0) \cdot d_1(a_1) \cdot R(w', a_0, a_1)(v') \cdot u(v') = \tilde{r}$$

and d_0 and d_1 are the optimal distribution choices for both players. We show that $\tilde{r} \leq r'$. Suppose by contradiction that $\tilde{r} > r'$. Then, let $\epsilon = \frac{\tilde{r} - r'}{2}$ and consider the function $f(v'') = \max(0, u(v'') - \epsilon)$. Clearly, (w', r', f) is a configuration in $\text{turn}(G)$. However, as π is a winning strategy from p the choice $\pi(p \cdot (w', r', f))$ contradicts the definition of u . So $\tilde{r} \leq r'$.

By assumption $t' < 1 - r'$. Let $\epsilon = 1 - r' - t'$. Consider now the function $f' : \text{succ}(w') \rightarrow [0, 1]$ such that $f'(v'') = 1 - u(v'') - \frac{\epsilon}{2}$. The minimax value of f' in $\text{turn}(\text{dual}(G))$ is at least $1 - \tilde{r} - \frac{\epsilon}{2} \geq 1 - r' - \frac{\epsilon}{2} > t'$. Hence, (w', t', f') is a configuration in $\text{turn}(\text{dual}(G))$.

We extend Player 0's strategy in $\text{turn}(\text{dual}(G))$ by choosing configuration (w', t', f') . Then, Player 1 answers by choosing a successor $(w \cdot v', f'(v'))$. Notice that it cannot be the case that $u(v') = 1$. Indeed, in such a case $f'(v')$ would be 0. So the path $\text{dual}(p)$ is extended by (w', t', f') and then $(w' \cdot v', f'(v'))$.

We now turn our attention to extension of the path p . By the choice of u , there is a function f such that (w', r', f) is a configuration in $\text{turn}(G)$, $\pi(p \cdot (w', r', f))$ is $(w' \cdot v', f(v'))$, and either $f(v') = u(v')$ or $f(v') < u(v') + \frac{\epsilon}{4}$. So we make Player 0 choose in $\text{turn}(G)$ the successor configuration (w', r', f) . Then, Player 1's winning strategy π instructs her to choose $(w' \cdot v', f(v'))$.

It follows that $f(v') \geq s(G, w' \cdot v')$. Otherwise, Player 0 has a winning strategy from $(w' \cdot v', f(v'))$ in contradiction with Player 1's strategy π being winning. Furthermore, π is winning from $p \cdot (w', r', f) \cdot (w' \cdot v', f(v'))$. Finally, as $f'(v') = 1 - u(v') - \frac{\epsilon}{2}$ and $f(v') < u(v') + \frac{\epsilon}{4}$ we conclude that $f'(v') < 1 - f(v')$.

- Suppose that $O(w') = > r''$ in G . Then $O(w') = \geq 1 - r''$ in $\text{dual}(G)$. It follows that (w', r') is a Player 0 configuration in $\text{turn}(G)$ and (w', t') is a Player 1 configuration in $\text{turn}(\text{dual}(G))$. Suppose that Player 1 chooses the next configuration (w', t'') in $\text{turn}(\text{dual}(G))$.

Now, this is similar to the previous case, as we have to continue from the configurations (w', r'') in $\text{turn}(G)$ and (w', t'') in $\text{turn}(\text{dual}(G))$ that are both Player 0 configurations and $t'' < 1 - r''$.

- Suppose that $O(w') = \geq r''$ in G . Then $O(w') = > 1 - r''$ in $\text{dual}(G)$. It follows that (w', r'') is a Player 1 configuration in $\text{turn}(G)$ and $(w', 1 - r'')$ is a Player 0 configuration in $\text{turn}(\text{dual}(G))$.

The winning strategy π instructs Player 1 to choose configuration (w', r''') such that $r''' < r''$.

As before, this is similar to the first case, as we have to continue from the configurations (w', r''') in $\text{turn}(G)$ and $(w', 1 - r'')$ in $\text{turn}(\text{dual}(G))$ that are both Player 0 configurations and $1 - r'' < 1 - r'''$.

Consider the two infinite plays played in $\text{turn}(G)$ and $\text{turn}(\text{dual}(G))$. Clearly, when projecting the two plays on the configurations in $V^+ \times (0, 1]$ that appear in

them and then on the configurations in V^+ we get exactly the same play. By assumption π is a winning strategy for Player 0 in $\text{turn}(G)$. Hence, this projection is not in α implying that the strategy constructed for Player 0 in $\text{turn}(\text{dual}(G))$ is indeed winning. \square

Corollary 3 *For every obligation Blackwell game G and every prefix w such that $O(w) \neq \perp$, $\text{val}_0(G, w) \in \{0, 1\}$.*

Proof: Consider a configuration v such that $O(w) = \bowtie r$. Then, by definition, the game $\text{turn}(G)$ starting from configuration (w, r') does not depend on the value r' . It follows that either Player 0 wins from (w, r') for all $r' \in (0, 1]$ or Player 1 wins from (w, r') for all $r' \in (0, 1]$. It follows that either $\text{val}_0(G, w) = 1$ or $\text{val}_0(G, w) = 0$. \square

4 Markov Chains with Obligations

We show that for Markov chains the measure of an obligation objective can be defined directly on the Markov chain. This direct characterization is generalized later and is crucial for the algorithmic analysis for finite games with parity objectives. We introduce the notion of a *choice set*, a set of obligations that Player 0 can meet. We then show that the definition of a value through a choice set and the definition in Section 3 coincide.

Consider a Markov chain $M = (S, P, L, s^{\text{in}})$. Let $\mathcal{G} = \langle \alpha, O \rangle$ be an obligation, where $\alpha \subseteq S^\omega$ is a Borel set of infinite paths and $O : S \rightarrow (\{\geq, >\} \times [0, 1]) \cup \{\perp\}$ is the obligation function. We can think about such a Markov chain as an obligation Blackwell game where A_0 and A_1 are singletons. Formally, $G_M = (S, \{a\}, \{a\}, R, \mathcal{G})$, where $R(s, a, a) = P(s)$ for all $s \in S$. As before, we are interested in sequences of locations, which correspond to prefixes of plays in G_M . Thus, we refer to them as prefixes also here. Let \hat{S} denote the set of prefixes $s^{\text{in}} \cdot S^*$. Let \mathcal{O} denote the set of locations $s \in S$ such that $O(s) \neq \perp$ and $\hat{\mathcal{O}}$ prefixes $w \cdot s \in \hat{S}$ such that $s \in \mathcal{O}$. That is, \mathcal{O} is the set of locations with a non-empty obligation and the set $\hat{\mathcal{O}}$ is the set of prefixes that end in a location in \mathcal{O} . We denote by $O(w)$ the obligation $O(s)$, where $w = w' \cdot s$. Let $\mathcal{N} = S \setminus \mathcal{O}$ denote the set of locations that have no obligation and $\hat{\mathcal{N}}$ denote the set of prefixes $\hat{S} \setminus \hat{\mathcal{O}}$. For a prefix w a *choice set* is $C_w \subseteq \hat{\mathcal{O}} \cap (\{w\} \cdot S^+)$. That is, it is a set of extensions of w that have obligations. For a prefix $w' \in \hat{S}$ and a choice set C_w , an infinite path $w' \cdot y$ is *good* if either (a) $y = x \cdot z$, $x \in \mathcal{N}^* \cdot \mathcal{O}$, and $w' \cdot x \in C_w$, or (b) $y \in \mathcal{N}^\omega$ and $w' \cdot y \in \alpha$. That is, either the first visit to \mathcal{O} after w' is in C_w or \mathcal{O} is never visited and the infinite path is in α . Let $\beta_{C_w}^{w'}$ denote the set of good paths of w' with choice set C_w . Given a choice set C_w and a prefix w' , the measure of \mathcal{G} from w' according to C_w is:

$$\text{Msr}_M^{\mathcal{G}}(w', C_w) = \frac{\text{Prob}_M(\beta_{C_w}^{w'})}{\text{Prob}_M(\{w'\} \cdot S^\omega)}.$$

A choice set C_w is *good* if the following two conditions hold:

- Every infinite path $\pi = s_0, s_1, \dots$ in M such that π has infinitely many prefixes in C_w is in α .
 - For every sequence $w' \in C_w$ we have $\text{Msr}_M^{\mathcal{G}}(w', C_w) \bowtie r$, where $O(w') = \bowtie r$.
- Let \mathcal{C}_w denote the set of good choice sets for w .

Consider a Markov chain $M = (S, P, L, s^{\text{in}})$ and an obligation $\mathcal{G} = \langle \alpha, O \rangle$. For prefix w the pre-value of w is

$$\tilde{v}(M, \mathcal{G}, w) = \sup_{C \in \mathcal{C}_w} \text{Msr}_M^{\mathcal{G}}(w, C).$$

Finally, we define the value of w . For a prefix w such that $O(w) \neq \perp$ we define $v(M, \mathcal{G}, w)$ to be 1 if $\tilde{v}(M, \mathcal{G}, w) \bowtie r$, where $O(w) = \bowtie r$, and $v(M, \mathcal{G}, w)$ is 0 otherwise. For a prefix w such that $O(w) = \perp$ we define $v(M, \mathcal{G}, w)$ to be $\tilde{v}(M, \mathcal{G}, w)$.

We note that in a choice set C , if there is some prefix $w \notin C$ such that $w \in \hat{O}$ then for every extension $w \cdot y$ of w there is no point in including $w \cdot y$ in C . Indeed, once a certain obligation is not included in C all the obligations that extend it are not important. We restrict attention to choice sets that satisfy this restriction.

We show that for every Markov chain and for every prefix the above definition of value coincides with definition through Martin-like reduction.

Theorem 5 *For every Markov chain M , obligation $\mathcal{G} = \langle \alpha, O \rangle$, and prefix $w \in S^*$ we have $v(M, \mathcal{G}, w) = \text{val}_0(G_M, w)$.*

The proof of Theorem 5 requires a refined analysis of a winning strategy in the uncountable game $\text{turn}(M)$ obtained from a Markov chain M . Using this analysis we extract a witness choice set in M from a winning strategy in the uncountable game.

Proof: We show that $\text{val}_0(G_M, w) \geq v(G, w)$.

- Fix $\epsilon > 0$. We have to show that if there is a choice set C that shows the value $v(G, w) - \epsilon$ then Player 0 can get the value $v(G, w) - \epsilon$ in $\text{turn}(G_M)$. This proof uses heavily Martin's proof of determinacy of Blackwell games [20]. Fix a Markov chain $M = (S, P, L, s^{\text{in}})$ for the rest of this proof.

Given a Borel winning set $\beta \subseteq S^\omega$, Martin defines a turn-based game \hat{G}_m that is slightly different to ours. Formally, $\hat{G}_m = \langle (S^+ \times [0, 1]) \cup (S^+ \times F), E, \hat{\beta} \rangle$, where F is the set of functions $\{f : S \rightarrow [0, 1]\}$ and

$$E = \{((w, v), (w, f)) \mid \sum_{s' \in S} P(s, s') \cdot f(s') \geq v\} \cup \{((w, f), (w \cdot s, v)) \mid f(s) \geq v\}.$$

For a set $P \subseteq S^+$, let $P \uparrow^w$ denote $P \cap \{w\} \cdot S^+$, i.e., exactly all suffixes of w in P . Then, based on determinacy of \hat{G}_m and measurability of β (since β is Borel), Martin's proof shows that $\frac{\text{Prob}_M(\beta \uparrow^w)}{\text{Prob}_M(S^+ \uparrow^w)} \geq v$ iff Player 0 wins \hat{G}_m from every configuration (w, v') for $v' < v$. That is, Player 0 announces the values she can derive from successors of w and Player 1 chooses a successor from which to show the value. Finally, $\hat{\beta}$ is the set of plays whose projection on S^+ has limit

in β . The strategy of Player 0 forces all infinite plays to be in $\hat{\beta}$.¹

We use Martin's result to show that whenever $\text{Msr}_M^G(w, C) \geq v$ then Player 0 wins in $\text{turn}(G_M)$ from (w, v') for every $v' < v$. Let $\delta = v - v'$. As $\text{Msr}_M^G(w, C) \geq v$, then according to Martin's proof Player 0 wins in \hat{G}_m from $(s, v' + \frac{\delta}{2})$. We use Player 0's strategy in \hat{G}_m to win in $\text{turn}(G_M)$. As we play, we maintain the requirement in $\text{turn}(G_m)$ always below the requirement in \hat{G}_m by repeatedly dividing the gap between the values in the two games by 2. It follows that in the i th round of playing the two games, the gap between the values is $\frac{\delta}{2^i}$. Furthermore, as Player 0's strategy in \hat{G}_m is winning it cannot be the case that the play created passes through an obligation prefix that is not in C (indeed, all continuations from this point are losing in \hat{G}_m). If on the other hand, a play passes through an obligation point that is in C , then the correspondence between the game $\text{turn}(G_M)$ and a new instance of \hat{G}_m from the new obligation point is created. Consider an obligation prefix w' such that $O(w') = \geq v'$. By goodness of C , $\text{Msr}_M^G(w', C) \geq v'$. In the game $\text{turn}(G_M)$ Player 1 moves to a configuration (w', v'', ϵ) , where $v'' < v'$. Thus, we can use the same argument and use Martin's game \hat{G}_m to continue the strategy in $\text{turn}(G_M)$. Consider an obligation prefix w' such that $O(w') = > v'$. By goodness of C , $\text{Msr}_M^G(w', C) > v'$. Hence, there is some v'' such that $v' < v'' < \text{Msr}_M^G(w', C)$ that can be used in Martin's game. Finally, consider an infinite play in $\text{turn}(G_M)$. If the play visits C infinitely often, then by C 's goodness, it is winning for Player 0. If the play visits C finitely often, according to Martin's result, the corresponding play is winning in \hat{G}_m implying that the play is in α .

In the other direction we show that $v(G, w) \geq \text{val}_0(G_M, w)$.

- In the other direction, a winning strategy for Player 0 in the game $\text{turn}(G_M)$ from (w, v) induces a good choice set C . We start by fixing the winning strategy σ of Player 0. For the sake of this proof we assume that Player 0 always plays by this strategy σ . Furthermore, from a prefix $w \in \hat{O}$ such that $O(w) = \geq v$, we know that Player 1 can choose every successor (w, v', ϵ) for $v' < v$. We restrict Player 1's choices to those v' such that $v' \geq 0$ and $v' = v - \frac{1}{n}$ for some $n \in \mathbb{N}$. Thus, when we say configuration is *reachable* we mean under these choices of Player 0 according to σ and for choices of Player 1 restricted in \geq -obligation configurations as explained.

The definition of the choice set C is quite involved as it has to take into account the infinitely many different strategies that are involved in showing an obligation of the form $\geq r$ (corresponding to each of the choices $v' = v - \frac{1}{n}$). We assume that the initial prefix w is an obligation such that $O(w) = \geq v$ for the value that interests us v . Indeed, this forces Player 0 to be able to win the game (w, v', ϵ) for every $v' = v - \frac{1}{n}$ for every $n \in \mathbb{N}$. The choice set we construct has to factor in these infinitely many different strategies. However, the same occurs whenever another obligation w' that extends w is reached for which $O(w') = \geq v''$ for

¹ The main difference between the two games (except for no obligations in Martin's version) is as follows. In our game the value promised by Player 0 is always slightly below the real value. Accordingly, we require that the weighted sum of values of the successors be strictly larger than the promised value. In Martin's version the weighted sum of values of successors may be equivalent to the promised values (or larger).

some v'' . It follows, that the choice set we construct must include the same construction for every obligation of the form $\geq v''$ that is encountered in the game. More formally, we have the following.

Assume that we start from prefix (w, r') such that $O(w) = \geq v$. Let T denote the set of prefixes reachable from w excluding w itself. For every $w' = w \cdot s_1 \cdots s_n \in T$, let $level(w')$ denote the number of locations s_i such that $s_i \in \mathcal{O}$ and $O(s_i) = \geq r'$ for $1 \leq i \leq n$ and $r' \in (0, 1]$. That is, $level(w')$ is the number of \geq -obligations on the way from w to w' excluding w itself but including w' (if appropriate). For every $i \geq 1$, let $T_i \subseteq T$ denote the set $T_i = \{w' \mid level(w') = i\}$. It follows that $T = \bigcup_{i \geq 1} T_i$ and for every i and j we have $T_i \cap T_j = \emptyset$. We now restrict attention (by induction) to a subset of the obligation configurations that appear in T .

Consider an obligation w' in T_1 . For every such obligations there is a minimal $n \in \mathbb{N}$ such that w' is reachable from $(w, v - \frac{1}{n}, \epsilon)$. We call this the *rank* of w , denoted $rank(w)$. Consider an obligation w' in T_i . Let s_1, \dots, s_i be the \geq -obligations on the way from w to w' and let v_1, \dots, v_i be the values of these obligations. Then, there is a minimal according to the lexicographic order (n_1, \dots, n_i) such that w' is reachable from w by Player 1 taking the choice $(s_j, v_j - \frac{1}{n_j}, \epsilon)$ from (s_j, v'_j) for the appropriate v'_j . As before, we call this the *rank* of w' , denoted $rank(w')$. We say that w' is *good* if for every \geq -obligation w'' on the path from w to w' we have that $rank(w'')$ is a prefix of $rank(w')$. That is, whenever w' is reachable through multiple choices of Player 1, we consider only the strategy Player 1 used from w' for the minimal choice of concession given on *all* \geq -obligations on the way to w' . We say that \geq -obligation prefix w'' is *good* if it appears in T and all \geq -obligation prefixes appearing on the path from w to w'' are good. That is, if it appears as part of one of the same “minimal” strategies. Let $C = \{w \in T \cap \hat{\mathcal{O}} \mid w \text{ is good}\}$. We note that the definition of C does not depend on G_M being derived from a Markov chain. Indeed the same definition is used in the proof of Theorem 7.

In order to show that C is a good choice set we have to prove two things. First, that every path that visits infinitely many obligations in C is in α . Second, that all obligations in C are met.

For the first claim we note that an infinite sequence of prefixes in C appears also in $\text{turn}(G_M)$ and from σ being a winning strategy must be in α . We have to show that the obligation of every $w' \in C$ is met. However, for this we can use again Martin’s reduction. For every prefix $w \in C$, the strategy of Player 0 in $\text{turn}(G_M)$ can be used to show a win in \hat{G}_m from (w, v') , where $O(w) = \bowtie r$ and either $\bowtie = \geq$ and $v' \geq r$ or $\bowtie = >$ and $v' > r$. Essentially, Player 0’s strategy in $\text{turn}(G_M)$ promises values for each prefix visited in a play. These values are larger than the values needed in \hat{G}_m and can be used to construct a strategy in \hat{G}_m . If some prefix in \mathcal{O} is reached, then clearly the play must be included in C as there is a strategy of Player 1 that makes it reachable. Furthermore, every infinite play that includes infinitely many prefixes in C can be forced by Player 1 in $\text{turn}(G_M)$ showing that it is in α . It follows that in the Markov chain we have $\text{Msr}_M^G(w, C) \bowtie r$.




Figure 5: An obligation Markov chain with no global good choice set.

□

The definition above uses the supremum over all good choice sets. We show that there is a choice set that attains the supremum.

Theorem 6 *For every Markov chain M , obligation $\mathcal{G} = \langle \alpha, O \rangle$, and prefix $w \in \widehat{S}$ there is a choice set $C \in \mathcal{C}$ such that $\text{Msr}_M^{\mathcal{G}}(w, C) = s(M, w)$.*

Proof: Fix a prefix w . There are choice sets $\{C_i\}_{i \in \mathbb{N}}$ such that $\text{Msr}_M^{\mathcal{G}}(w, C_i) \geq s(G, w) - \frac{1}{2^i}$.

Let $C'_0 = C_0$. Consider a set C_{i+1} . Let

$$C'_{i+1} = C_{i+1} \setminus \{w' \in C_{i+1} \mid \exists w'' \in C'_j \text{ for } j \leq i \text{ s.t. } w' = w'' \cdot y \text{ for some } y\}.$$

We set $C = \bigcup_{i=1}^{\infty} C'_i$. We show that C is a good choice set.

Consider an infinite path that visits infinitely many prefixes in C . Clearly, all of the prefixes in C belong to the same set C_i for some $i \in \mathbb{N}$. Hence, by C_i being a good choice set, the path is in α .

Consider a point $w' \in C$. Let i be the minimal such that $w' \in C'_i$. Then, for every extension $w'' = w' \cdot y$ such that $w'' \cdot y \in C_i$ we have $w'' \in C'_i$. Indeed, as w' does not have a prefix in C'_j for all $j < i$, so is the case for w'' . Then, as $\text{Msr}_M^{\mathcal{G}}(w', C_i)$ satisfies the obligation of w' then so does $\text{Msr}_M^{\mathcal{G}}(w', C)$.

We can show that $\text{Msr}_M^{\mathcal{G}}(w, C) = s(G, w)$. Indeed, if it were smaller than $s(G, w)$ then there is an i such that $\text{Msr}_M^{\mathcal{G}}(w, C_i) > \text{Msr}_M^{\mathcal{G}}(w, C)$. It must be the case that C_i includes extensions of w that are not in C , contradicting the definition of C . □

We can show that in some cases there is no one good choice set that covers all possible prefixes. Consider for example the Markov chain in Figure 5. Suppose that the path in which s_1 appears infinitely often is not in α . Clearly, for every prefix $p = s_1 \cdots s_1$ the pre-value of this configuration is 1. Indeed, the choice set that includes exactly $p \cdot s_1$ proves that. However, this choice set, establishes the value of $p \cdot s_1$ as $\frac{2}{3}$, which is, as required, more than $\frac{1}{3}$. A choice set that shows the values of *all* prefixes simultaneously, has to include all prefixes $s_1 \cdots s_1$. Thus, the infinite path $s_1 \cdot s_1 \cdots$ is visited infinitely often by this choice set and it cannot be good.

5 Finite Turn-Based Stochastic Parity Games with Obligations

We extend the results from obligation Markov chains to finite turn-based stochastic parity games with obligations, and show that the value function in such games has an




Figure 6: Finite turn-based stochastic parity game with obligations requiring memory. Diamonds are Player 1 configurations and circles are stochastic configurations. Priorities in the range $[0..1]$ next to state define a parity acceptance condition. Only configurations v_6 and v_8 have priority 0.

alternate direct characterization using choice sets. The direct characterization is crucial to present algorithms to solve finite turn-based stochastic parity games with obligations. The simpler definition does not generalize to infinite games, Blackwell games, or more general winning conditions. In these more complicated games optimal strategies do not always exist. For such games, we need a more elaborate construction that captures the winning with ϵ -optimal strategies (and non-existence of optimal strategies), as done in Section 3.

We reuse the notation $G = ((V, E), (V_0, V_1, V_p), \kappa, \mathcal{G})$ for *turn-based stochastic parity games with obligations*. Here, $\mathcal{G} = \langle \alpha, O \rangle$ is a goal and α is derived from a parity condition $c : V \rightarrow [0..k]$. Strategies are defined as before. Given a prefix $w \in V^+$ and two strategies σ and π , we denote by $w(\sigma, \pi)$ the Markov chain obtained from G by using the strategies σ and π starting from prefix w . Then, we define the value of Player 0 in the prefix w in the game to be $v(G, w) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} v(w(\sigma, \pi), \mathcal{G}, w)$.

We note that if we extend the definition of a measure of a choice set so that bad choice sets give measure 0 for all configurations then the following holds:

$$\begin{aligned} v(G, w) &= \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \sup_{C_w} \text{Msr}_{w(\sigma, \pi)}^{\mathcal{G}}(w, C_w) \geq \\ &\geq \sup_{\sigma \in \Sigma} \sup_{C_w} \inf_{\pi \in \Pi} \text{Msr}_{w(\sigma, \pi)}^{\mathcal{G}}(w, C_w) \end{aligned}$$

This follows from properties of supremum and infimum. In the proof of Theorem 7 we actually show that $v(G, w) \leq \text{val}_0(G, w) \leq \sup \sup \inf(\dots)$. Hence, the two are actually equivalent. Formally, we show that for every finite turn-based stochastic parity game with obligations and for every prefix the two values $v(G, w)$ and $\text{val}_0(G, w)$ coincide.

Theorem 7 *For all finite turn-based stochastic obligation parity games G and prefix $w \in V^+$, $v(G, w) = \text{val}_0(G, w)$.*

As another illustration of the notion of a choice set consider the game in Figure 6. In order to use v_1 , Player 0 has to win more than $\frac{1}{2}$ from that configuration. Choosing the self loop from v_3 to itself or the edge from v_3 to v_1 only makes things worse (though, each can be chosen a finite number of times). So the only option from v_3 is to go back to v_2 so that the probability of getting from v_2 to v_4 is 1. If from v_4 ,

Player 0 chooses to go to v_7 the value is $\frac{1}{2}$ which does not satisfy the obligation of v_1 . Going from v_4 to v_5 , on the other hand, and upon returning from v_5 to v_4 proceeding to v_7 fulfills all obligations. Indeed, the value for v_1 is 1 as all paths eventually reach v_5 , and the value for v_5 is $\frac{3}{4}$ as the loop to itself through v_6 is winning and the paths from v_5 to v_4 and then on to v_7 have value $\frac{1}{4}$. It follows that a possible choice set for this game is $C = \{v_1, v_1v_2(v_3v_2)^*v_4v_5(v_6^+v_5)^*\}$. Indeed, Player 0 has a strategy reaching from v_1 to $v_1v_2(v_3v_2)^*v_4v_5$ with probability 1. She has a strategy from $v_1v_2(v_3v_2)^i v_4v_5(v_6^jv_5)^k$ to either reach $v_1v_2(v_3v_2)^i v_4v_5(v_6^jv_5)^k v_6^+v_5$ or win the parity objectives with probability $\frac{3}{4}$. We note that Player 0 uses its first visit to v_4 to go to v_5 , in order to boost the probability needed for the obligation of v_1 , and in subsequent visits goes to v_7 .

We show that for a finite turn-based stochastic parity game G , from a winning strategy in the uncountable game $\text{turn}(G)$, we can construct a regular witness choice set C' . Using the regular witness and the fact that for finite turn-based stochastic ω -regular games finite-memory optimal strategies exist we obtain Theorem 7.

Proof: As before, let $\mathcal{O} = \{v \in V \mid O(v) \neq \perp\}$, $\mathcal{N} = \{v \in V \mid O(v) = \perp\}$. Let $\mathcal{O}_i = \{v \in \mathcal{O} \mid c(v) = i\}$ and $\mathcal{N}_i = \{v \in \mathcal{N} \mid c(v) = i\}$ be the obligation and non-obligation configurations with priority i . Similarly, let $\mathcal{O}_{\geq i} = \{v \in \mathcal{O} \mid c(v) \geq i\}$ and $\mathcal{N}_{\geq i} = \{v \in \mathcal{O} \mid c(v) \geq i\}$.

We show that $\text{val}(G, w) \geq \text{val}_0(G, w)$.

- Consider a winning strategy for Player 0 in $\text{turn}(G)$ that starts in (w, r) . We can extract from it a set C of obligations that are used. This is done just like in the proof of Theorem 5. Clearly, every path that visits infinitely many prefixes in C is in α as it appears also in $\text{turn}(G)$ as before. We modify C to a set C' that we can show is good. We construct C' by induction on the number of obligations passed on the way from (w, r) .

Consider a prefix $w' \in C \cup \{w\}$. Let $ob(w')$ be the set of obligations $w'' \in C$ such that $w'' = w' \cdot w'''$ and $w''' \in \mathcal{N}^* \cdot \mathcal{O}$. That is, $ob(w')$ is the set of obligations directly reachable from w' without passing through other obligations. Furthermore, annotate every prefix w'' in $ob(w')$ by the minimal priority occurring in w''' , where $w'' = w' \cdot w'''$ and $w''' \in \mathcal{N}^* \cdot \mathcal{O}$. For every prefix $w' \in C \cup \{w\}$ we define a set $obs(w') \subseteq V \times [0..k]$, where $[0..k]$ are the priorities of the parity condition. Formally, $obs(w')$ is the set of pairs (v', i') such that some $w'' \in ob(w')$ is annotated by i' and the last configuration in w'' is v' .

Let \mathcal{N}_i and \mathcal{O}_i denote the configurations in G whose priority is i . Let $\mathcal{N}_{\geq i}$ and $\mathcal{O}_{\geq i}$ denote the configurations in G whose priority is at least i .

We now construct C' by induction. We label every prefix $p' \in C'$ by a prefix $p \in C$ that is the *reason* for inclusion of p' in C' . Consider the configuration w . By construction, for every $(v, i) \in obs(w)$ there is a prefix $w_{(v,i)} \in C$ such that $w_{(v,i)} \in ob(w)$ and $w_{(v,i)}$ ends in v . We add to C' all the prefixes $w \cdot p$, where p is in the following set (restricted to those reachable from w):

$$\bigcup_{(v,i) \in obs(w)} (\mathcal{N}_{\geq i}^* \cdot \mathcal{N}_i \cdot \mathcal{N}_{\geq i}^* \cdot \{v\} \cup \mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_i \cap \{v\}))$$

Furthermore, every prefix $w \cdot p$ for $p \in (\mathcal{N}_{\geq i}^* \cdot \mathcal{N}_i \cdot \mathcal{N}_{\geq i}^* \cdot \{v\} \cup \mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_i \cap \{v\}))$ is

labeled by $w \cdot w_{(v,i)}$. We now continue by induction. Consider a prefix $p' \in C'$ that is labeled by prefix $p \in C$. By construction, for every $(v, i) \in \text{obs}(p)$ there is a prefix $p_{(v,i)} \in C$ such that $p_{(v,i)} \in \text{ob}(p)$ and $p_{(v,i)}$ ends in v . We add to C' all the prefixes $p' \cdot p''$, where p'' is in the following set (restricted to those reachable from p'):

$$\bigcup_{(v,i) \in \text{obs}(p)} (\mathcal{N}_{\geq i}^* \cdot \mathcal{N}_i \cdot \mathcal{N}_{\geq i}^* \cdot \{v\} \cup \mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_i \cap \{v\}))$$

Furthermore, every prefix $p' \cdot p''$ for $p'' \in (\mathcal{N}_{\geq i}^* \cdot \mathcal{N}_i \cdot \mathcal{N}_{\geq i}^* \cdot \{v\} \cup \mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_i \cap \{v\}))$ is labeled by $p \cdot p_{(v,i)}$.

This completes the construction of C' . We have to show that C' is a good choice set. That is, every infinite path in G that visits infinitely many configurations in C' is fair and the strategy of Player 0 in G establishes all the obligations posed by C' .

The fact that every infinite path in G that visits infinitely many prefixes in C' is fair can be deduced by following the labels in C of prefixes in C' . Consider such an infinite sequence of prefixes p_0, p_1, \dots in C' and their respective labels w_0, w_1, \dots from C . By construction, the minimal priority visited in the extension of p_i to p_{i+1} is the minimal priority visited in the extension of w_i to w_{i+1} . Furthermore, the sequence w_0, w_1, \dots corresponds to a path in G that visits infinitely many configurations in C . As C is obtained from $\text{turn}(G)$, it follows that the limit of w_0, w_1, \dots is fair. That is, the limit of p_0, p_1, \dots satisfies the parity objective. We conclude that the limit of p_0, p_1, \dots is fair as well.

We now have to show that all obligations in C' are met. Consider a prefix $p' \in C'$ labeled by prefix $p \in C$. Both p' and p end in the same configuration $v \in V$ of G . It follows that the obligation $O(v)$ is fulfilled in $\text{turn}(G)$. Assume that $O(v) = \geq r$. It follows that Player 0 wins in $\text{turn}(G)$ from (p, r') for every $r' < r$. Recall the sets $\text{ob}(p) \subseteq C$ and $\text{obs}(p) \subseteq V \times [0..k]$. The winning in $\text{turn}(G)$ from (p, r') for every $r' < r$ can be translated to a win in \hat{G} for the goal $\text{ob}(p) \cdot V^\omega$ for every $r'' < r$, where \hat{G} is the game obtained from G by Martin's reduction. Thus, the value of $\text{ob}(p) \cdot V^\omega \cup (\alpha \cap \mathcal{N}^\omega)$ in G is r . We now consider the following goal γ in G :

$$\gamma = \bigcup_{(v,i) \in \text{obs}(p)} ((\mathcal{N}_{\geq i}^* \cdot \mathcal{N}_i \cdot \mathcal{N}_{\geq i}^* \cdot \{v\}) \cup (\mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_i \cap \{v\}))) \cup (\alpha \cap \mathcal{N}^\omega)$$

In particular, γ contains at least all the extensions p' such that $p \cdot p' \in \text{ob}(p)$ as well as $\alpha \cap \mathcal{N}^\omega$. Furthermore, γ can be translated to a parity goal in G by including a simple monitor for the minimal parity encountered along the path. As $\text{reach}(\text{ob}(p)) \cup (\alpha \cap \mathcal{N}^\omega) \subseteq \gamma$ it follows that the value of γ in G is at least r . However, values in finite turn-based stochastic parity games are attained. That is, there is a strategy for Player 0 such that the value of γ according to this strategy is at least r . It follows that by using this strategy Player 0 can ensure the obligation of p in G . The case that $O(v) = > r$ is simpler, as Player 0 wins directly from (v, r) in $\text{turn}(G)$.

It follows that C' is a good choice set and that Player 0 has a strategy that ensures that all obligations in C' are met.

In the other direction we show that $v(G, w) \leq \text{val}_0(G, w)$.

- Suppose by way of contradiction that $r = v(G, w) > \text{val}_0(G, w)$. Let $t < r$ be such that $t > \text{val}_0(G, w)$. By definition of $\text{val}_0(G, w)$, Player 1 wins in $\text{turn}(\text{dual}(G))$ from $(w, 1 - t)$. This winning strategy induces a good choice set T just like in the previous proofs. Note that this set is good for Player 1. Thus, every path that visits infinitely many configurations in T *does not* satisfy the acceptance condition. As in the other direction of the proof, the set T can be extended to a good choice set T' such that Player 1 has a strategy to enforce all obligations in T' (in the dual game). We show that T' proves that the value of w in G cannot be r . Fix a strategy σ of Player 0 in G . We show that the strategies of Player 1 in G that enforce the obligations in T' (in the dual game) induce a strategy $\pi \in \Pi$ such that every choice set in $G_{\sigma, \pi}$ that shows the value r cannot be good. Notice, that the set \mathcal{O} is the same in G and $\text{dual}(G)$. Hence, there is a strategy π for Player 1 in G that achieves the value $1 - t$ for Player 1. Consider the strategy σ and assume that it achieves the value r for Player 0 in G for some choice set C . Then, as $1 - t + r > 1$, and $\mathcal{N}^\omega \cap \alpha$ and $\mathcal{N}^\omega \cap \bar{\alpha}$ are disjoint, it follows that C and T' have a non-empty intersection such that the strategy σ reaches $C \cap T'$.

We now proceed by induction. Consider a configuration $w' \in C \cap T'$. There are two cases, either $O(w') = \geq r'$ or $O(w') = > r'$.

- Suppose that $O(w') = \geq r'$. In this case, the obligation of w' in $\text{dual}(G)$ is $> 1 - r'$. It follows that there is some value $r'' < r'$ such that the value of $(T' \cdot V^\omega) \cup (\mathcal{N}^\omega \cap \bar{\alpha})$ in G for Player 1 is $1 - r''$. Then, as $1 - r'' + r' > 1$, and $\mathcal{N}^\omega \cap \alpha$ and $\mathcal{N}^\omega \cap \bar{\alpha}$ are disjoint, it follows that C and T' have a non-empty intersection such that the strategy σ reaches $C \cap T'$.
- Suppose that $O(w') = > r'$. In this case, there is a value $r'' > r$ such that σ must attain the goal $(C \cdot V^\omega) \cup (\mathcal{N}^\omega \cap \alpha)$ with probability r'' . At the same time Player 1 can force the goal $(T' \cdot V^\omega) \cup (\mathcal{N}^\omega \cap \bar{\alpha})$ with probability $1 - r'$. As $r'' + 1 - r' > 1$ the sets C and T' have a non-empty intersection such that the strategy σ reaches $C \cap T'$.

Continuing by induction we create a path that visits infinitely many configurations in T' and in C . It follows that C cannot be a good choice set. \square

Corollary 4 *For every finite turn-based stochastic parity game with obligations G and prefix $w \in V^+$, there are strategies $\sigma \in \Sigma$ and $\pi \in \Pi$ such that $v(G, w) = v(w(\sigma, \pi), \mathcal{G}, w)$. Furthermore, for every strategy $\sigma' \in \Sigma$ and $\pi' \in \Pi$ we have $v(w(\sigma', \pi'), \mathcal{G}, w) \leq v(G, w) \leq v(w(\sigma, \pi'), \mathcal{G}, w)$.*

Proof: This follows from the proofs of Theorems 5 and 7. Consider a configuration w . Suppose that $v(G, w) = r$. Then, for every n there is a strategy σ_n such that for every $\pi \in \Pi$ the value $v(G_{\sigma_n, \pi}(w), \mathcal{G}, w) \geq r - \frac{1}{n}$. Furthermore, there is a good choice set C_n such that the goal $(C_n \cdot V^\omega) \cup (\mathcal{N}^\omega \cap \alpha)$ is enforced with probability at least $r - \frac{1}{n}$. As in the proof of Theorem 5 the different choice sets $\{C_n\}_{n>0}$ can be combined to a single choice set C . Furthermore, the choice set C has a simple structure as in the proof of Theorem 7. It follows that Player 0 can enforce the goal

$(C \cdot V^\omega) \cup (\mathcal{N}^\omega \cap \alpha)$ with probability larger than $r - \frac{1}{n}$ for every n . As G is finite it must be that $(C \cdot V^\omega) \cup (\mathcal{N}^\omega \cap \alpha)$ can be enforced with probability r .

The proof that Player 1 also has an optimal strategy is similar. \square

6 Algorithmic Analysis of Obligation Games

We give algorithms for solving obligation Blackwell games in two cases. First, in case that in every path in the game, the number of transitions between an obligation configuration and a non-obligation configuration is bounded. In this case, we show that obligation Blackwell games can be reduced to a sequence of turn-based stochastic games. Second, in case that the game is finite and the winning condition is a parity condition. In this case, we give an exponential time algorithm for computing the value of the game.

6.1 Reduction to Stochastic Games

Essentially, this is the solution adopted in [14] for solving acceptance of uniform p-automata. We partition the game to regions where there are no transitions between obligation configurations and non-obligation configurations. A region that consists only of non-obligation configurations can be thought of as a stochastic game. A region that consists only of obligation configurations can be thought of as a turn-based (non-stochastic) game. More formally, we have the following.

Consider an obligation Blackwell game $G = (V, A_0, A_1, R, \mathcal{G})$, where $\mathcal{G} = \langle \alpha, O \rangle$. We say that a configuration v is *pure* if for every $a_0 \in A_0$ and $a_1 \in A_1$ we have $R(v, a_0, a_1)$ is pure. We say that the game is *uniform* if all the following holds.

- There is a partition $\{V_i\}_{i \in \mathbb{N}}$ of V such that for every i we have, either (i) for every $v \in V_i$ we have $O(v) = \perp$ or (ii) for every $v \in V_i$ we have $O(v) \neq \perp$ or v is pure.
- We say that $V_i \leq V_{i'}$ if there are some $v \in V_i$, $v' \in V_{i'}$, $a_0 \in A_0$, and $a_1 \in A_1$ such that $R(v, a_0, a_1)(v') > 0$. The partition must also satisfy that every chain according to \leq is finite.

Theorem 8 *The computation of the value of a uniform obligation Blackwell game G can be reduced to the solution of multiple Blackwell games.*

Proof: Let $\{V_i\}_{i \in \mathbb{N}}$ be the partition of the game G . By assumption, consider a set V_i such that there is no other set $V_{i'}$ such that $V_i < V_{i'}$. Consider a prefix $w = w' \cdot v$ such that $v \in V_i$. Clearly, the extension of this prefix to a play in G remains forever in V_i .

Suppose that for all $v \in V_i$ we have $O(v) = \perp$. Let $G' = (V^*, A_0, A_1, R, \alpha)$ be the game obtained from G by restricting attention to configurations reachable from w . The game G' is a normal Blackwell game and hence the value of every configuration in $\{w\} \cdot V_i^*$ is well defined.

Suppose that for all $v \in V_i$ we have $O(v) \neq \perp$ or v is pure. Consider the turn-based game $G' = ((V^* \cup V^* \times 2^V, E), (V^*, V^* \times 2^V), \alpha')$, where we restrict V^* to configurations reachable from w and E and α' are as follows. Consider a prefix $u' \cdot v'$

and a set $S \subseteq V$. If $O(v') = \perp$, we say that S is *possible* from $u' \cdot v'$ if there is $d_0 \in \mathcal{D}(A_0)$ such that for all $d_1 \in \mathcal{D}(A_1)$ we have

$$\sum_{a_0 \in A_0} \sum_{a_1 \in A_1} \sum_{v'' \in S} R(v', a_0, a_1)(v'') \bowtie p, \quad (1)$$

where $O(v') = \bowtie p$. If v' is pure, we say that S is *possible* from $u' \cdot v'$ if there is $a_0 \in A_0$ such that for all $a_1 \in A_1$ the unique configuration v'' such that $R(v, a_0, a_1)(v'') = 1$ is in S . Notice, that this is like considering a pure configuration as having the obligation ≥ 1 .

- $E = \{(w' \cdot v', (w' \cdot v', S)) \mid S \text{ possible from } w' \cdot v'\} \cup \{((w', S), (w' \cdot v')) \mid v' \in V\}$.
- α' includes all infinite paths such that the limit of their projection on V^* is in α . This is in effect equivalent to the reduction to $\text{turn}(G)$ when restricted to $\{w\} \cdot V_i^*$.

Consider now a set V_i and a configuration $w = w' \cdot v$ such that $v \in V_i$. Suppose, by induction, that for all configurations $u \cdot v'$ such that $v' \in V_{i'}$ for $V_i < V_{i'}$ a value has already been computed.

If for every $v \in V_i$ we have $O(v) = \perp$, then a similar reduction to a normal Blackwell game by plugging in the value of precomputed configurations gives the value of all configurations in $\{w\} \cdot V_i^*$.

If for every $v \in V_i$ we have $O(v) \neq \perp$, then a similar reduction to a turn-based game can be done. This time, value of precomputed configurations has to be combined in the small minimax games as in Equation 1. \square

We note that this is a “meta”-algorithm. Consider a uniform obligation Blackwell game G and the partition V_1, \dots, V_n showing that it is uniform. Suppose that every V_i reduces to a Blackwell game that can be analyzed algorithmically. Then, from Theorem 8, the game G can be analyzed algorithmically.

6.2 Finite Turn-based Stochastic Obligation Parity Games

We show that values in finite turn-based stochastic parity games with obligations (POG, for short) can be computed in exponential time and decision problems regarding values lie in $\text{NP} \cap \text{co-NP}$.

We give a nondeterministic algorithm for finding a maximal (wrt to inclusion) choice set, which calls the computation of values in stochastic parity games as a subroutine. Then, the value of a configuration in the game can be computed by computing the value of reaching the choice set computed by the algorithm or winning the parity condition without reaching other obligations. By results of previous sections, dualization of the game gives the value of the opponent. It follows that the decision regarding the value is also in co-NP.

We now give an algorithm that decides and computes values in G . A *dependency* for $v \in \mathcal{O}$ is either $C_v = \perp$ or $C_v \subseteq (\mathcal{O} \times [0..k])$. That is, C_v is either *undefined* or a (possibly empty) set of pairs of obligation configurations annotated by priorities. A *game dependency* is a set $\{C_v\}_{v \in \mathcal{O}}$. A game dependency is *good* if the following conditions hold:




Figure 7: Illustration of a dependency. Priorities in the range $[0..4]$ next to state define a parity acceptance condition.

1. If for some $v \in \mathcal{O}$ we have $(v', i) \in C_v$ then $C_{v'} \neq \perp$.
2. For every infinite sequence $(v_0, i_0), (v_1, i_1), \dots$ such that for every j we have $(v_{j+1}, i_{j+1}) \in C_{v_j}$ the minimal priority occurring infinitely often in i_0, i_1, \dots is even.
3. For every $v \in \mathcal{O}$ such that $C_v \neq \perp$ we have $\text{val}_0(G', v) \bowtie r$, where $O(v) = \bowtie r$ and G' is the game G considered as a turn-based stochastic game with the goal γ :

$$\bigcup_{(v', i) \in C_v} \left(\begin{array}{c} (\mathcal{N}_{\geq i}^* \cdot \mathcal{N}_i \cdot \mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_{\geq i} \cap \{v'\}) \cdot V^\omega) \cup \\ (\mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_i \cap \{v'\}) \cdot V^\omega) \cup \\ (\alpha \cap \mathcal{N}^\omega) \end{array} \right)$$

Informally, for an obligation v with non-empty dependency, the dependency indeed shows that the obligation is met: Player 0 can force (i) winning the original winning condition while never reaching another obligation or (ii) reaching an obligation v' that v depends on, with i , the required parity, being the minimal visited along the way.

We illustrate the notion of a dependency using Figure 7. The obligation of s_1 is $\frac{3}{4}$. The probability to reach s_1 from itself is 1. However, the paths $(s_1 s_2^+ s_4 s_6)^\omega$ have a minimal priority of 1 and are losing. It follows that the only winning paths are $(s_1 s_2^+ s_4 s_5)^\omega$, $(s_1 s_3^+ s_4 s_5)^\omega$, and $(s_1 s_3^+ s_4 s_6)^\omega$. Thus, s_1 depends on reaching s_1 with minimal priority 0 (through s_3) and on reaching s_1 with minimal priority 2 (through s_2 and s_5). This satisfies the three conditions as (1) s_1 has a defined dependency, (2) every cycle visits either the minimal priority 0 or 2, and (3) the probability of reaching s_1 with minimal priority 0 is $\frac{1}{2}$, the probability of reaching s_1 with minimal priority 2 is $\frac{1}{4}$, and the probability of not reaching s_1 is 0. So the total probability is $\frac{3}{4}$, which fulfils the obligation of s_1 . Adding an obligation of $\geq \frac{1}{2}$ at s_4 , changes the dependency. Now, s_1 depends on reaching s_4 with priority 0 or 2 and s_4 depends on reaching s_1 with priority 3. However, if the obligation of s_4 is set to $> \frac{1}{2}$, then there is no good dependency. Indeed, this would mean that *whenever* s_4 is reached the path through s_6 must be included. Then, the path from s_1 through s_2 can not be part of the dependency as this would create a cycle with minimum priority 1 and the obligation of s_1 is no longer fulfilled. The dependency for the game in Figure 6 is v_1 depends on reaching v_5 with priority 1 and v_5 depends on reaching itself with priority 0. This is a good dependency as (a) the only cycle in it is v_5 reaching itself with priority 0 (b) from v_1 Player 0 has a strategy that ensures that v_5 is reached with probability 1, and (c) from v_5 Player 0 has a strategy that ensures that either v_5 is reached with minimal priority

0 encountered or getting to v_8 and staying there (with no obligations on the way) with probability $\frac{3}{4}$.

The nondeterministic algorithm is as follows. We guess a game dependency $\{C_v\}_{v \in \mathcal{O}}$. The size of $\{C_v\}_{v \in \mathcal{O}}$ is polynomial in $|V|$. We check that $\{C_v\}_{v \in \mathcal{O}}$ is good by doing the following. First, checking that if $(v', i) \in C_v$ then $C_{v'} \neq \perp$ can be completed in polynomial time by scanning all the sets C_v . Second, checking that all cycles induced by C_v have a minimal even parity in them can be completed in polynomial time by drawing the graph of connections between the different configurations in \mathcal{O} for which $C_v \neq \perp$ and searching for a cycle with minimal odd priority. Third, ensuring that the values in the different turn-based stochastic games fulfill the obligations can be achieved in $\text{NP} \cap \text{co-NP}$ by Theorem 3. Finally, consider the goal γ' :

$$\gamma' = (\mathcal{N}^* \cdot \{v : V_c \neq \perp\} \cdot V^\omega) \cup (\alpha \cap \mathcal{N}^\omega)$$

We evaluate whether $\text{val}_0(G, w) \bowtie r$ by checking whether $\text{val}_0(G', w) \bowtie r$, where G' is the turn-based stochastic game obtained from G by considering the goal γ' . This can be checked in $\text{NP} \cap \text{co-NP}$. To compute the value $\text{val}_0(G, w)$ we compute the value of w in G' . This can be computed in exponential time. Notice that the values of γ' in G' correspond to the value $s(G, w)$ and not $\text{val}_0(G, w)$. For obligation configurations we must compare the result with the required obligations. If the obligation is met, the value $\text{val}_0(G, w)$ is 1. Otherwise, it is 0. Overall, if *all* the nondeterministic guesses are made up-front (i.e., the dependency *and* the winning strategies in *all* games) then the global size of the witness is polynomial and all the checks can be completed in polynomial time. Overall, the decision problem is in $\text{NP} \cap \text{co-NP}$, and the values can be computed in exponential time.

We apply the algorithm on the example in Figure 7. As analyzed above, the dependency for s_1 is $(s_1, 0)$ and $(s_1, 2)$. This proves that the value *for all configurations* is 1. Indeed, for every configuration in the game the probability of reaching s_1 at least once is 1. Once s_1 is reached for the first time, the more complex reliance on the choice set that is extracted from the dependency is required. Applying the algorithm on the game in Figure 6 we see that the value of v_1, v_2, v_3, v_4 , and v_6 is 1 as from them Player 0 can reach v_5 with probability 1. The pre-value of v_5 is $\frac{3}{4}$ as it reaches itself with probability $\frac{1}{2}$ and wins the parity condition (reaching v_8) with probability $\frac{1}{4}$. As this matches its obligation its value is 1.

Algorithm correctness follows from the following Lemmas.

Lemma 3 *There is a memoryless winning strategy in $\text{turn}(G)$.*

Proof: According to the proof of Theorem 7 from every obligation used as part of the winning strategy in $\text{turn}(G)$, there is a simple goal that leads to the next frontier of used obligations. Namely, given the sets $ob(p)$ and $obs(p)$ the goal is:

$$\gamma = \bigcup_{(v,i) \in obs(p)} ((\mathcal{N}_{\geq i}^* \cdot \mathcal{N}_i \cdot \mathcal{N}_{\geq i}^* \cdot \{v\}) \cup (\mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_i \cap \{v\}))) \cup (\alpha \cap \mathcal{N}^\omega)$$

The structure of this goal implies that there is a strategy with memory linear in the number of priorities in the game that achieves an optimal value for this goal. However,

an obligation configuration $v \in \mathcal{O}$ may appear infinitely often in $\text{turn}(G)$, each time using a different strategy.

Consider now all the possible strategies for Player 0 in G with a goal γ as above with memory bounded by the number of priorities. Clearly, the number of such strategies is finite. In particular, for every obligation configuration $v \in \mathcal{O}$ there is a finite number of strategies that are used in $\text{turn}(G)$. We now construct a finite parity game G' based on these strategies. For every prefix $p \cdot v$ such that $v \in \mathcal{O}$ used in $\text{turn}(G)$ add a Player 0 configuration v to G' . For every strategy σ that is used from the prefix $p \cdot v$ in $\text{turn}(G)$ we add a Player 1 configuration σ to G' . For every pair (v', i) such that $v' \in \mathcal{O}$ and i is a priority such that application of σ from $p \cdot v$ reaches a configuration v' with priority i being the minimal visited along the way we add the Player 0 configuration (σ, v', i) to G' . We add edges to G' as follows. From configuration v we add edges to all strategies σ used from $p \cdot v$ for some p . From strategy σ we add edges to all triplets (σ, v', i) . From configuration (σ, v', i) we add an edge to v' . We set the priority of (σ, v', i) to be i and priorities of all other configuration to be the maximal possible priority.

The game G' is a finite parity game and we know that Player 0 wins G' based on the combination of the winning strategies in $\text{turn}(G)$. It follows from Theorem 3 that there is a memoryless winning strategy for Player 0 in G' . However, a memoryless winning strategy in G' induces a unique choice of a strategy from every obligation configuration in $\text{turn}(G)$ leading to a memoryless winning strategy in $\text{turn}(G)$. \square

Lemma 4 *An obligation configuration v fulfills $\text{val}_0(G, v) = 1$ iff there is a good game dependency $\{C_v\}_{v \in \mathcal{O}}$ such that $C_v \neq \emptyset$.*

Proof: The existence of a good game dependency clearly shows that the obligation of v can be met.

In the other direction, if the obligation of v can be met, this means that Player 0 wins in $\text{turn}(G)$ from (v, r') for every $r' \in (0, 1]$. Furthermore, a choice set of a very particular form can be extracted as in the proof of Theorem 7. According to Lemma 3 Player 0 has a memoryless winning strategy in $\text{turn}(G)$. We note further, that in $\text{turn}(G)$, if an obligation configuration (w, r') occurs, the game below (w, r') does not depend on the value r' . Thus, if two obligations $w \cdot v$ and $w' \cdot v$ and the configurations $(w \cdot v, r)$ and $(w' \cdot v, r')$ occur in $\text{turn}(G)$ the extension of the game below both is identical. It follows, that the memoryless strategy behaves exactly the same from all obligations $w \cdot v$ and $w' \cdot v$ for the same obligation configuration v . Then $\text{obs}(w \cdot v) = \text{obs}(w' \cdot v)$ for all $w, w' \in V^*$. So for an obligation v appearing in $\text{turn}(G)$ we can use the set $\text{obs}(w \cdot v)$ for some prefix w as the dependency C_v . For every obligation v' not appearing in $\text{turn}(G)$ we set $C_{v'} = \perp$. We have to show that this induces a good game dependency. First, if we have $(v', i) \in C_v$ then it follows that some prefix $w' \cdot v'$ is reachable from a prefix $w \cdot v$. Thus, C'_v must be defined. Second, every cycle in $\{C_v\}_{v \in \mathcal{O}}$ with minimal odd priority corresponds to an infinite path in the good choice set with a minimal odd priority, which is impossible. Third, it must be the case that $\text{val}_0(G', v) \bowtie r$, where G' is obtained from G by considering the goal γ :

$$\gamma = \bigcup_{(v', i) \in C_v} ((\mathcal{N}_{\geq i}^* \cdot \mathcal{N}_i \cdot \mathcal{N}_{\geq i}^* \cdot \{v'\} \cdot V^\omega) \cup (\mathcal{N}_{\geq i}^* \cdot (\mathcal{O}_i \cap \{v\}) \cdot V^\omega)) \cup (\alpha \cap \mathcal{N}^\omega)$$

Indeed, this is the exact construction of the choice set from $\text{turn}(G)$ as in the proof of Theorem 7, where it is proven that it is also good. \square

Lemma 5 *For every configuration v , $\text{val}_0(G, v) = r$ iff there is a good game dependency $\{C_{v'}\}_{v' \in \mathcal{O}}$ such that $\text{val}(G', v) = r$, where G' is obtained from G by considering the goal γ .*

$$\gamma = (\mathcal{N}^* \cdot \{v : V_c \neq \perp\} \cdot V^\omega) \cup (\alpha \cap \mathcal{N}^\omega)$$

Proof: As before, if $\text{val}(G', v) = r$ then clearly $\text{val}_0(G, v) \geq r$. In the other direction, we consider all the obligations appearing in the choice set showing that $\text{val}_0(G, v) = r$. According to the previous lemma, these obligations require a good game dependency. Finally, the value $\text{val}_0(G, v)$ is exactly the reachability of the good choice set or winning parity without reaching obligations. \square

Theorem 9 *For a POG G and a prefix $w \in V^+$, the values $\text{val}_0(G, w)$ and $\text{val}_1(G, w)$ can be computed in exponential time and whether $\text{val}_0(G, w) \bowtie r$ can be decided in $NP \cap co-NP$.*

7 p-Automata

In [14], we defined uniform p-automata and showed that they are a complete abstraction framework for pCTL. Acceptance of Markov chains by uniform p-automata was defined through a cumbersome and complicated reduction to a series of turn-based stochastic parity games. Here, using obligation games, we give a clean definition of acceptance by p-automata. What's more, obligation games allow us to define acceptance by general p-automata and remove the restriction of uniformity. To simplify presentation we remove the notion of $*$ -transitions (see [14]).

We assume familiarity with basic notions of trees and (alternating) tree automata. For set T , let $B^+(T)$ be the set of positive Boolean *formulas* generated from elements $t \in T$, constants tt and ff , and disjunctions and conjunctions:

$$\varphi, \psi ::= t \mid \text{tt} \mid \text{ff} \mid \varphi \vee \psi \mid \varphi \wedge \psi \quad (2)$$

Formulas in $B^+(T)$ are finite even if T is not.

For set Q , the set of states of a p-automaton, we define *term* sets $\llbracket Q \rrbracket_>$ as follows.

$$\llbracket Q \rrbracket_> = \{\llbracket q \rrbracket_{\bowtie p} \mid q \in Q, \bowtie \in \{\geq, >\}, p \in [0, 1]\}$$

Intuitively, a state $q \in Q$ of a p-automaton and its transition structure model a probabilistic path set. So $\llbracket q \rrbracket_{\bowtie p}$ holds in location s if the measure of paths that begin in s and satisfy q is $\bowtie p$.

An element of $Q \cup \llbracket Q \rrbracket_>$ is therefore either a state of the p-automaton, or a term of the form $\llbracket q \rrbracket_{\bowtie p}$. Given $\varphi \in B^+(Q \cup \llbracket Q \rrbracket_>)$, its *closure* $\text{cl}(\varphi)$ is the set of all subformulas of φ . For a set Φ of formulas, let $\text{cl}(\Phi) = \bigcup_{\varphi \in \Phi} \text{cl}(\varphi)$.

Definition 1 A p -automaton A is a tuple $\langle \Sigma, Q, \delta, \varphi^{\text{in}}, \alpha \rangle$, where Σ is a finite input alphabet, Q a set of states (not necessarily finite), $\delta: Q \times \Sigma \rightarrow B^+(Q \cup \llbracket Q \rrbracket_{>})$ the transition function, $\varphi^{\text{in}} \in B^+(\llbracket Q \rrbracket_{>})$ the initial condition, and α a parity acceptance condition.

In general, p -automata have *states*, Markov chains have *locations*, and games *configurations*.

For every \mathbb{AP} , p -automata $A = \langle 2^{\mathbb{AP}}, Q, \delta, \varphi^{\text{in}}, \alpha \rangle$ have $\text{MC}_{\mathbb{AP}}$ as set of inputs. For $M = (S, P, L, s^{\text{in}}) \in \text{MC}_{\mathbb{AP}}$, we define whether A accepts M by a reduction to a turn-based stochastic parity game with obligations. The language of A is $\mathcal{L}(A) = \{M \in \text{MC}_{\mathbb{AP}} \mid A \text{ accepts } M\}$.

We construct a game $G_{M,A} = ((V, E), (V_0, V_1, V_p), \kappa, \mathcal{G})$. A configuration of $G_{M,A}$ corresponds to a subformula appearing in the transition of A and a location in M . Configurations with a term of the form $\llbracket q \rrbracket_{\bowtie p}$ correspond to obligations. All other configurations have no obligations. The Markov chain is accepted if the configuration $(\varphi^{\text{in}}, s^{\text{in}})$ has value 1 in $G_{M,A}$.

Formally, we define $G_{M,A}$ as follows. Let $G_{M,A} = ((V, E), (V_0, V_1, V_p), \kappa, \mathcal{G})$, where the components of $G_{M,A}$ are as follows.

- $V = S \times \text{cl}(\delta(Q, \Sigma))$.
- $V_0 = \{(s, \psi_1 \vee \psi_2) \mid s \in S \text{ and } \psi_1 \vee \psi_2 \in \text{cl}(\delta(Q, \Sigma))\}$.
- $V_1 = \{(s, \psi_1 \wedge \psi_2) \mid s \in S \text{ and } \psi_1 \wedge \psi_2 \in \text{cl}(\delta(Q, \Sigma))\}$.
- $V_p = S \times (Q \cup \llbracket Q \rrbracket_{>})$.
- The set of edges E is defined as follows.

$$\begin{aligned} E = & \{((s, \varphi_1 \wedge \varphi_2), (s, \varphi_i)) \mid i \in \{1, 2\}\} \cup \\ & \{((s, \varphi_1 \vee \varphi_2), (s, \varphi_i)) \mid i \in \{1, 2\}\} \cup \\ & \{((s, q), (s', \delta(q, L(s)))) \mid s' \in \text{succ}(s)\} \cup \\ & \{((s, \llbracket q \rrbracket_{\bowtie p}), (s', \delta(q, L(s)))) \mid s' \in \text{succ}(s)\} \end{aligned}$$

- $\kappa((s, q), (s', \delta(q, L(s)))) = \kappa((s, \llbracket q \rrbracket_{\bowtie p}), (s', \delta(q, L(s)))) = P(s, s')$.
- $\mathcal{G} = \langle \tilde{\alpha}, O \rangle$, where
 - For $q \in Q$ and $p \in [0, 1]$ we have $\tilde{\alpha}(s, q) = \alpha(q)$, $\tilde{\alpha}(s, \llbracket q \rrbracket_{\bowtie p}) = \alpha(q)$. For every other configuration c we set $\tilde{\alpha}(c)$ to the maximal possible priority.
 - For $q \in Q$ and $p \in [0, 1]$ we have $O(s, \llbracket q \rrbracket_{\bowtie p}) = \bowtie p$. For every other configuration c , we have $O(c) = \perp$.

As obligation games are well defined it follows that it is well defined whether a p -automaton accepts a Markov chain.

Theorem 10 Given a finite p -automaton A and a finite Markov chain M , we can decide whether $M \in \mathcal{L}(A)$ in time exponential in the number of states of A and locations of M .

Proof: This follows from the polynomial construction of the finite-state turn-based stochastic obligation parity game $G_{M,A}$ for the Markov chain M and p -automata A , and Theorem 9. \square

The definition in [14] restricts attention to uniform p -automata. Such automata restrict the cycles in the transition graph of p -automata. We recall the definition of uniform p -automata. In doing so, we differentiate states q' appearing within a term in $\llbracket Q \rrbracket_>$ (bounded transition) from q' appearing “free” in the transition of a state q (unbounded transition). In this way, a p -automaton $A = \langle \Sigma, Q, \delta, \dots \rangle$ determines a labeled, directed graph $G_A = \langle Q', E, E_b, E_u \rangle$:

$$\begin{aligned} Q' &= Q \cup \text{cl}(\delta(Q, \Sigma)) \\ E &= \{(\varphi_1 \wedge \varphi_2, \varphi_i), (\varphi_1 \vee \varphi_2, \varphi_i) \mid \varphi_i \in Q' \setminus Q, \\ &\quad i \in \{1, 2\}\} \cup \{(q, \delta(q, \sigma)) \mid q \in Q, \sigma \in \Sigma\} \\ E_u &= \{(\varphi \wedge q, q), (q \wedge \varphi, q), (\varphi \vee q, q), (q \vee \varphi, q) \mid \\ &\quad \varphi \in Q', q \in Q\} \\ E_b &= \{(\llbracket q \rrbracket_{\bowtie p}, q) \mid \llbracket q \rrbracket_{\bowtie p} \in \llbracket Q \rrbracket_>\} \end{aligned}$$

Elements $(\varphi, q) \in E_u$ are *unbounded* transitions; elements $(\varphi, q) \in E_b$ are *bounded* transitions; and elements of E are called *simple* transitions. Note that E , E_u , and E_b are pairwise disjoint. Let $\varphi \preceq_A \tilde{\varphi}$ iff there is a finite path from φ to $\tilde{\varphi}$ in $E \cup E_b \cup E_u$. Let \equiv be $\preceq_A \cap \preceq_A^{-1}$ and $(\llbracket \varphi \rrbracket)$ the equivalence class of φ with respect to \equiv . Each $(\llbracket \varphi \rrbracket)$ is an SCC in the directed graph G_A .

Definition 2 [14] A p -automaton A is called *uniform* if: (a) For each cycle in G_A , its set of transitions is either in $E \cup E_b$ or in $E \cup E_u$. (b) There are only finitely many equivalence classes $(\llbracket \varphi \rrbracket)$ with $\varphi \in Q \cup \text{cl}(\delta(Q, \Sigma))$.

That is, A is uniform, if the full subgraph of every equivalence class in \preceq_A contains only one type of non-simple transitions. Also, all states $q' \in Q$ or formulas φ occurring in $\delta(q, \sigma)$ for some $q \in Q$ and $\sigma \in \Sigma$ can be classified as unbounded, bounded, or simple – according to SCC $(\llbracket q \rrbracket)$. Intuitively, the cycles in the structure of a uniform p -automaton A take either no bounded edges or no unbounded edges. Uniformity allowed to define acceptance for p -automata through the solution of a sequence of stochastic games.

Theorem 11 p -automata (Definition 1) extend the definition of uniform p -automata (Definition 2).

Proof: The obligation Blackwell game resulting from the composition of a uniform p -automaton with a Markov chain is a uniform turn-based obligation game. From Theorem 8 it follows that its value is obtained from solving a sequence of turn-based stochastic games and turn-based games. This gives rise to exactly the definition of acceptance through a sequence of turn-based stochastic games and turn-based games as in [14]. \square

Closure under union and intersection is easy due to alternation (see [14] for details). Closure under complementation follows from our determinacy result for obligation games.

8 Conclusions and Future Work

We introduced obligations, a structural winning condition that complements winning conditions on paths. We show that Blackwell games with Borel objectives and obligations are well defined. We then present a simpler definition of value for Markov chains with Borel objectives and obligations and for finite turn-based stochastic parity games with obligations. Based on the simpler definition we give algorithms for analyzing finite turn-based stochastic parity games with obligations. We then use games with obligations to define acceptance by unrestricted p-automata, showing that the new definition generalizes a previous definition for uniform p-automata.

This is one of the rare cases in games that arise in verification that determinacy of games does not immediately follow from Martin's result that Blackwell games with Borel objectives are determined. The proof of determinacy uses elements from Martin's determinacy proof but introduces new concepts that were not needed in that proof. These new concepts are required due to the more elaborate nature of games with obligations.

Our work gives rise to many interesting questions. For example, determining the complexity of other types of games such as Streett, Rabin, Muller, and quantitative games with obligations.

Finally, many questions regarding the theory of p-automata remain open. For instance, understanding the different transition modes of such automata (i.e., alternation vs. nondeterminism vs. determinism) and conversions between the different modes. A related question is that of feasibility of algorithmic questions such as emptiness of p-automata, which generalizes the satisfiability problem of pCTL.

References

- [1] A. Aziz, V. Singhal, F. Balarin, R. Brayton, and A. Sangiovanni-Vincentelli. It usually works: the temporal logic of stochastic systems. In *Proc 7th Int. Conf. on Computer Aided Verification*, volume 939 of *Lecture Notes in Computer Science*, pages 155–165. Springer, 1995.
- [2] P. Billingsley. *Probability and Measure*. Wiley, 2008.
- [3] K. Chatterjee. *Stochastic ω -regular Games*. PhD thesis, University of California at Berkeley, 2007.
- [4] K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Quantitative stochastic parity games. In *Symposium on Discrete Algorithms*, pages 114–123. SIAM, 2004.
- [5] F. Ciesinski and C. Baier. Liquor: A tool for qualitative and quantitative linear time analysis of reactive systems. In *Proc. 3rd International Conference on Quantitative Evaluation of Systems*, pages 131–132. IEEE Computer Society, 2006.
- [6] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.

- [7] E.A. Emerson and C.-L. Lei. Modalities for model checking: Branching time logic strikes back. In *Proc. 12th ACM Symp. on Principles of Programming Languages*, pages 84–96, 1985.
- [8] E. Grädel, W. Thomas, and T. Wilke. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
- [9] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.
- [10] S. Hart and M. Sharir. Probabilistic propositional temporal logics. *Information and Control*, 70(2–3):97–155, 1986.
- [11] S. Hart, M. Sharir, and A. Pnueli. Termination of probabilistic concurrent programs. In *Proc. 9th ACM Symp. on Principles of Programming Languages*, pages 1–6, 1982.
- [12] T.A. Henzinger, O. Kupferman, and S. Rajamani. Fair simulation. *Information and Computation*, 173(1):64–81, 2002.
- [13] A. Hinton, M.Z. Kwiatkowska, G. Norman, and D. Parker. Prism: A tool for automatic verification of probabilistic systems. In *Proc. 12th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 3920 of *Lecture Notes in Computer Science*, pages 441–444. Springer, 2006.
- [14] M. Huth, N. Piterman, and D. Wagner. Weak p-automata: Acceptors of Markov chains. In *Proc. 7th International Conference on Quantitative Evaluation of Systems*. IEEE Computer Society Press, 2010.
- [15] M. Jurdziński. Deciding the winner in parity games is in $\text{up} \cap \text{co-up}$. *Information Processing Letters*, 68(3):119–124, 1998.
- [16] O. Kupferman, M.Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *Journal of the ACM*, 47(2):312–360, 2000.
- [17] K. G. Laresn and B. Jonsson. Specification and refinement of probabilistic processes. In *Proc. 6th IEEE Symp. on Logic in Computer Science*, pages 266–277. IEEE Computer Society Press, 1991.
- [18] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
- [19] D.A. Martin. Borel determinacy. *Annals of Mathematics*, 65:363–371, 1975.
- [20] D.A. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- [21] M. Mio. *Game Semantics for Probabilistic μ -Calculi*. PhD thesis, University of Edinburgh, 2012.

- [22] M. Mio. Probabilistic modal μ -calculus with independent product. *Logical Methods in Computer Science*, 8(4), 2012.
- [23] M. Mio and A. Simpson. Łukasiewicz μ -calculus. In *Proc. Fixpoints in Computer Science*, 2013.
- [24] J. Von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.
- [25] H.L. Royden and P. Fitzpatrick. *Real Analysis*. Prentice Hall, 2010.
- [26] W. Thomas. Languages, automata, and logic. *Handbook of Formal Language Theory*, III:389–455, 1997.
- [27] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. 26th IEEE Symp. on Foundations of Computer Science*, pages 327–338, 1985.